

Individual consequences of occupational decline

Per-Anders Edin

Tiernan Evans

Georg Graetz

Sofia Hernnäs

Guy Michaels

The Institute for Evaluation of Labour Market and Education Policy (IFAU) is a research institute under the Swedish Ministry of Employment, situated in Uppsala.

IFAU's objective is to promote, support and carry out scientific evaluations. The assignment includes: the effects of labour market and educational policies, studies of the functioning of the labour market and the labour market effects of social insurance policies. IFAU shall also disseminate its results so that they become accessible to different interested parties in Sweden and abroad.

Papers published in the Working Paper Series should, according to the IFAU policy, have been discussed at seminars held at IFAU and at least one other academic forum, and have been read by one external and one internal referee. They need not, however, have undergone the standard scrutiny for publication in a scientific journal. The purpose of the Working Paper Series is to provide a factual basis for public policy and the public policy discussion.

More information about IFAU and the institute's publications can be found on the website www.ifau.se

ISSN 1651-1166

Individual Consequences of Occupational Decline^a

Per-Anders Edin^b, Tiernan Evans^c, Georg Graetz^d, Sofia Hernnäs^e, & Guy Michaels^f

June 4, 2019

Abstract

What are the earnings and employment losses that workers suffer when demand for their occupations declines? To answer this question we combine forecasts on occupational employment changes, which allow us to identify unanticipated declines; administrative data on the population of Swedish workers, spanning several decades; and a highly detailed occupational classification. We find that, compared to similar workers, those facing occupational decline lost about 2-5 percent of mean cumulative earnings from 1986-2013. But workers at the bottom of their occupations' initial earnings distributions suffered considerably larger losses. These earnings losses are partly accounted for by reduced employment, and increased unemployment and retraining.

KEYWORDS: Technological change, Occupations, Inequality.

JEL CLASSIFICATION: O33, J24, J62.

^aCorresponding author: Guy Michaels (g.michaels@lse.ac.uk). Edin, Graetz, Hernnäs: Uppsala University. Evans, Michaels: LSE. We thank Wenkui Liu for excellent research assistance. We are indebted to Michael Böhm, Adrian Adermon, Magnus Gustavsson, and Linus Liljeberg for sharing code and data. For helpful comments and discussions we thank our discussants Jonathan Haskell and Robert Seamans, as well as Daron Acemoglu, David Autor, Abhijit Banerjee, Paul Beaudry, Matias Cortes, Nicole Fortin, Peter Fredriksson, Thomas Lemieux, Marco Manacorda, Alan Manning, Mattias Nordin, Barbara Petrongolo, Steve Pischke, Oskar Nordström Skans, Assaf Peretz, Jonathan Vogel, and Alwyn Young; participants at AEA meetings in Philadelphia and Atlanta, CEP Policy Roundtable, CEPR/ECB Conference, ECB Heads of Research workshop, Employment in Europe conference in Cyprus, ENEF conference in Sussex, Essex Workshop on Innovation and Labor Markets, Nordic Summer Institute in Labour Economics, Queen Mary Economic Workshop, Royal Society and British Academy workshop, and TU Munich workshop on Innovation, Technological Change, and International Trade; and seminar participants at IIES, IFN, Konstanz, Linnaeus University, LSE, Mannheim, MIT, Sveriges Riksbank, UBC, Uppsala University, and the World Bank. We are grateful to the Centre for Economic Performance and to Forte: Swedish Research Council for Health, Working Life and Welfare for generous financial support. Graetz thanks the faculty and staff of ifo Institute Munich, where part of the work on this paper was done, for their hospitality. Any opinions and errors are our own.

^b Per-Anders Edin, Uppsala University

^c Tiernan Evans, T.Evans2@lse.ac.uk, London School of Economics.

^d Georg Graetz, Georg.Graetz@nek.uu.se, Uppsala university.

^e Sofia Hernnäs, Sofia.Hernnas@nek.uu.se, Uppsala university

^f Guy Michaels, G.Michaels@lse.ac.uk, London School of Economics

1 Introduction

What are the long-run employment and earnings losses incurred by individual workers when demand for their occupation declines? This question lies at the heart of policy debates on responses to technologies that replace workers (Acemoglu and Restrepo, 2019), and is relevant for broader discussions on labor market transformations due to technological change (see for instance Brynjolfsson and McAfee, 2014, Autor, 2015, and Caselli and Manning, 2018). New labor-replacing technologies no longer threaten only machine operatives and clerical workers. Self-driving vehicles may reduce the employment of drivers (Campbell, 2018), and artificial intelligence software challenges professionals such as lawyers and financial investors (Susskind and Susskind, 2015) and even fashion designers (Scheiber, 2018). This is causing considerable angst. It is therefore important to understand how costly occupation-replacing technologies are for workers, since this informs our thinking about individual welfare, inequality, and human capital investments. It is also important for public policy decisions on taxation, redistribution, retirement, and education, and may even have broader political consequences (Marx, 1867; Caprettini and Voth, 2017; Dal Bo, Finan, Folke, Persson, and Rickne, 2019).

In this paper, we investigate the consequences of large, negative occupational demand changes for individual workers' careers. We combine forecasts on occupational employment changes, which allow us to identify unanticipated declines; population-level administrative data spanning several decades; and a highly detailed occupational classification. We are therefore able to study similar workers who perform similar tasks and have similar expectations of future occupational employment trajectories, but experience different actual employment changes.¹

To measure occupational decline, we use the US Occupational Outlook Handbook (Bureau of Labor Statistics, 1986, henceforth OOH), which allows us to identify which occupations declined in the US since the mid-1980s; to check whether occupational declines had likely technology drivers; and to gauge expectations of employment growth at the time. For reasons that we discuss below, our baseline definition of occupational decline requires that employment contracted by at least 25 percent since the mid-1980, though we also explore many alternative

¹For example, the employment of typists has nearly vanished, while secretaries' employment has grown.

definitions. We match the occupational information from the OOH to individual-level panel data on the entire Swedish population. Thus, we utilize the best aspects of both countries' data: the US data allow us to study occupational changes over time and separate unanticipated changes from anticipated ones, while the Swedish data let us follow individuals who differ in their exposure to occupational declines, but were otherwise very similar.

Focusing on cohorts that were in prime working age from the mid-1980s till the mid-2010s, we study how cumulative long-run outcomes (such as earnings and employment) differ for those who in 1985 worked in occupations that subsequently declined. We control for the initial sorting of workers into declining occupations by gender, age, education, income, and location in 1985. We show that conditional on these controls, those in occupations that subsequently declined had similar cognitive and non-cognitive skills and parental education and earnings, and similar pre-1985 earnings, as other workers. In some specifications we add other controls, including measures of occupation-varying life-cycle profiles and predictors of occupational change, which allow us to isolate unanticipated occupational declines, and we also control for broad occupation and industry dummies. We provide evidence that our various specifications yield credible bounds on the impact of occupational decline.

We also confirm that both our OOH-based measure of occupational decline and the predicted changes in US employment correlate strongly with the employment changes in Sweden. Specifically, Swedish workers who started out in occupations that subsequently declined were exposed to employment growth that was 20-40 log points lower than in non-declining occupations. We find that compared to workers with similar characteristics, those exposed to occupational decline lost about 5 percent of mean cumulative pre-tax earnings and 2 percent of mean cumulative employment. And compared to similar workers in similar occupations and industries, the cumulative earnings losses were only around 2 percent, and the cumulative employment losses were around 1 percent. We also find that those in declining occupations were significantly more likely to have exited their 1985 occupation by 2013. If occupational demand curves slope downward, this higher exit likely mitigated the earnings losses for those who remained in declining occupations.

While mean earnings losses from occupational decline were around 2-5 percent, those in the bottom tercile of their occupation's earnings distribution in 1985 suffered larger losses, amounting to 8-11 percent. Those at the bottom (and possibly also the top) of their occupation's earnings distribution were also less likely to remain in their starting occupation.

We also find that occupational decline increased the cumulative time spent in unemployment (accounting for roughly a third of lost employment) and retraining (accounting for just under ten percent of lost employment). Moreover, occupational decline led to slightly earlier retirement among middle-aged (in 1985) workers.

In addition, we show that our main findings are largely unchanged when we restrict the analysis to occupational declines that are explicitly linked to specific technological changes, as documented in the OOH.

Finally, we estimate mean earnings and employment losses from occupational decline that were similar in magnitude or possibly even smaller using micro data from the US (National Longitudinal Survey of Youth 1979); the NLSY estimates (at least for earnings) are, however, noisier than those using Swedish data. Nevertheless, these findings suggest that our estimates of losses from occupational decline may generalize to settings beyond Sweden.

To frame our empirical analysis of the consequences of occupational decline, we construct a Roy (1951) model with occupational demand shocks. As discussed above, we find that the largest earnings losses from occupational decline in Sweden are incurred by those who earned the least within their initial occupations. This finding is inconsistent with the frictionless Roy model, but it is consistent with a version where occupational switching costs decline in the workers' ability in the destination occupation. Moreover, our empirical analysis sheds light on the nature of the occupational switching costs, as almost half of the employment losses we estimate are accounted for by a combination of increased retraining and especially unemployment.

Our model can account for additional empirical findings when we also allow for worker displacement. In this case, those with lower initial within-occupation earnings rank suffer larger earnings losses as a result of occupational decline; switchers' earnings losses may be larger than those of stayers; and switching probabilities are U-shaped in initial earnings, whereby low-

earning workers switch occupations if displaced, while high-earning workers switch regardless of displacement when faced with occupational decline.

Occupational decline is a salient feature of the evolution of labor markets (Goldin, 2000). But despite its importance, we know relatively little about its consequences of occupational decline for individuals' careers. While there is a large literature on the costs of occupational mobility, we are not aware of previous estimates of workers' earnings losses from negative occupation-level demand changes.²

Our paper is distinct from panel studies of workers who differ in the routineness of their jobs.³ A key difference is that we can compare similar workers, even doing similar work, with different exposure to occupational decline.⁴ Our paper also differs conceptually from studies of mass layoffs. Occupational decline can sometimes be managed through retirement and reduced hiring, allowing workers to change jobs without leaving employment; and occupational decline need not entail severe spillovers for local labor markets, unlike mass layoffs (Gathmann, Helm, and Schönberg, 2018). While magnitude comparisons across studies should be interpreted with caution, the mean loss that we find from occupational decline is generally lower than the loss from mass layoffs.⁵ Finally, our paper also differs from studies of trade shocks, which affect import-competing firms and industries, while the changes we study typically affect individual workers within firms.

Our paper is also related to Dauth, Findeisen, Suedekum, and Woessner (2018), who explore how workers fare who are exposed to industrial robots; ours differs by exploring the consequences of a broader set of changes in occupational employment. Furthermore, our paper

²Cortes and Gallipoli (2017), Kambourov and Manovskii (2009), Pavan (2011), and Sullivan (2010) estimate the human capital losses associated with switching occupations. An older literature, including Neal (1995) and Parent (2000) studies the cost of moving across industries, while in other related work Gathmann and Schönberg (2010) and Poletaev and Robinson (2008) focus on task-specific human capital. Changes in the task content of existing occupations (for instance Spitz-Oener, 2006), while also potentially relevant, are outside the scope of our study due to data limitations.

³See for example Cortes (2016) and Autor and Dorn (2009).

⁴Also related is independent work by Schmillen (2018), who studies employment shocks faced by German apprentices, although our paper differs in its research question, econometric inference, and outcomes.

⁵Studies of mass layoffs in Sweden find losses of 4-6 percent of annual earnings in the 5-10 years following displacement (Eliason and Storrie, 2006; OECD, 2015). In the US losses from displacement are generally larger, and range from 7-14 percent of earnings (Davis and Von Wachter, 2011), or possibly even higher for workers who were highly attached to their firms (Jacobson, LaLonde, and Sullivan, 1993). Galaasen and Kostøl (2018) and Bana (2019) explore how mass layoffs' effects differ for occupations facing negative demand shocks, but their focus is still on mass layoffs

is related to the literature on possible future displacement due to technological changes. Forecasts of occupational displacement range from almost 50 percent (Frey and Osborne, 2017) to around 10 percent (Arntz, Gregory, and Zierahn, 2017, who obtain a lower estimate by taking into account within-occupation heterogeneity in tasks). At the same time, Bessen (2016) concludes that technology has, at least so far, not been a net destroyer of jobs. Even if this benign aggregate trend continues, however, some occupations may be replaced by technology, and our study offers a way to assess the losses from occupational displacement.

Sweden's economy and labor market institutions constitute the backdrop to most of our empirical analysis. During the period of our study, the Swedish economy experienced a deep recession in the early 1990s and a milder one in 2008 (Lindbeck, 1997; Gottfries, 2018), and we find that earnings losses in declining occupations were worse during those recessions. Wage inequality in Sweden increased during the 1980s and 1990s and remained relatively stable thereafter (Skans, Edin, and Holmlund, 2009). Swedish labor market institutions have been characterized by strong labor unions and substantial public spending on labor market policies. Unions have generally embraced technological changes to promote productivity and wage gains, while expecting that active labor market policy will help displaced workers find work (Edin and Holmlund, 1995). There is, indeed, some evidence that Sweden's occupational retraining programs raise earnings (Vikström and van den Berg, 2017), so they may have contributed to the modest losses from occupational decline that we find.⁶ At the same time, our finding of similarly modest mean earnings and employment losses from occupational decline in the US, suggests that workers find ways to mitigate losses from occupational decline even in other settings.

The remainder of our paper is organized as follows. Section 2 presents our model, Sections 3 and 4 discuss our data and empirical strategy, respectively, Section 5 presents our results, and Section 6 concludes.

⁶Another feature of Swedish labor market institutions are so-called employment security agreements reached between labor unions and business associations, and administered by works councils. These agreements stipulate counselling of laid-off workers to minimize the duration of their unemployment. We do not consider these agreements important in driving our results because, first, private sector blue-collar workers were only covered from 2004 onwards, and second, a careful evaluation of these agreements does not find strong support for positive treatment effects (Andersson, 2017).

2 Occupational decline in a Roy model

This section presents a simple model to help us frame our empirical investigation. We consider two occupations, one of which is hit by a negative demand shock. We investigate how workers' likelihood of leaving the affected occupation, and their earnings losses, depend on their initial earnings. Starting from a standard frictionless Roy (1951) model, we successively introduce positive and potentially heterogeneous costs of switching occupation; as well as the possibility that workers are displaced from their jobs and incur a cost to find a new job even when remaining in their initial occupation. Finally, we consider how workers' sorting differs when the negative demand shock is anticipated in advance. A complete, self-contained exposition of the model is given in the online appendix. Here we only summarize the main elements.

We consider a competitive economy with a continuum of individuals indexed by i who live for two periods $t \in \{1, 2\}$ and each supplies a unit of labor inelastically each period. There are two occupations indexed by $k \in \{A, B\}$ for the workers to choose from. Workers' per-period log earnings are given by $y_{ikt} = \pi_{kt} + \alpha_{ik} - c_{ikt}$, where π_{kt} is the time-varying and stochastic (log) price of a unit of output in occupation k , α_{ik} is the time-invariant (log) amount of output that worker i produces in occupation k , and $c_{ikt} \geq 0$ is a time cost related to occupational switching, which we discuss below.⁷ There are no saving opportunities and earnings are consumed immediately. We define the life-time expected utility function as $\mathbb{E}[y_{ik1} + \beta y_{ik2}]$, where $\beta > 0$ is a discount factor. In each period, workers choose the occupation that maximizes their expected utility. As a normalization, we assume that workers always choose occupation A if indifferent. Since we focus our analysis on relative wages, we define $\tilde{\pi}_t \equiv \pi_{Bt} - \pi_{At}$ and assume for simplicity that $\tilde{\pi}_1 = 0$.⁸ Prices are determined in equilibrium by supply and demand. However, here we take them as given, and analyze the consequences of a change to prices occurring in period 2 for occupational sorting and earnings. Note that the second period may be interpreted as all periods following this change, so β could be larger than one. For simplicity, we assume that

⁷The time cost may reflect search or retraining (or both); we assume throughout that a worker's wage equals the value of her marginal product, $e^{\pi_{kt} + \alpha_{ik}}$. We thus abstract from any job-level rents that may arise in the presence of search frictions.

⁸We do not claim to identify any aggregate gains from technological change, and we do not model them here.

α_{iA} and α_{iB} are independent and both uniformly distributed between zero and some finite but possibly large number $\bar{\alpha}$. We explain in the online appendix that our main results are robust to alternative assumptions about the joint skills distribution.

In period 2, there is a negative demand shock to occupation A such that $\pi_{A2} - \pi_{A1} = -d$ and $\tilde{\pi}_2 = d, d > 0$. This may be due to labor-replacing technology becoming available, or cheaper, in occupation A . We are interested in the consequences of the shock for the earnings of workers who start out in occupation A , under various assumptions about switching costs and anticipation of the price change. Formally, let $l_i \equiv \mathbb{E}[y_{i2} | \alpha_{iA}, k_{i1} = A, D_A = 0] - \mathbb{E}[y_{i2} | \alpha_{iA}, k_{i1} = A, D_A = 1]$ be the expected earnings loss in period 2 that results from the shock, conditional on worker i starting out in occupation A , and conditional on her ability (and hence earnings rank) α_{iA} , where the occurrence of the shock is indicated by $D_A \in \{0, 1\}$. Similarly, l_i^{switch} and l_i^{stay} denote the earnings losses further conditioned on leaving and staying, respectively, and p_i is the probability of switching.⁹ The overall loss is given by

$$l_i = l_i^{\text{stay}} - p_i \left(l_i^{\text{stay}} - l_i^{\text{switch}} \right). \quad (1)$$

As long as there is no displacement then $l_i^{\text{stay}} = d$ and by revealed preference $l_i^{\text{switch}} \leq d$, so that $l_i \leq d$. Thus, switching enables workers to mitigate the losses from occupational decline. In the online appendix we show that in each version of our model, $\frac{\partial p_i}{\partial d} \geq 0, \frac{\partial l_i}{\partial d} \geq 0$ (with strict inequalities for some i): the larger the drop in demand, the more workers switch, and the higher are earnings losses. Furthermore, $\frac{\partial l_i}{\partial \alpha_{iA}} = -\frac{\partial p_i}{\partial \alpha_{iA}} \left(l_i^{\text{stay}} - l_i^{\text{switch}} \right) + p_i \frac{\partial l_i^{\text{switch}}}{\partial \alpha_{iA}}$. In other words, losses decrease with initial within-occupation earnings rank if the switching probability is increasing and the loss of switchers decreasing in initial earnings rank, $\frac{\partial p_i}{\partial \alpha_{iA}} > 0$ and $\frac{\partial l_i^{\text{switch}}}{\partial \alpha_{iA}} < 0$.

⁹Formally,

$$\begin{aligned} l_i &\equiv l_i(\alpha_{iA}, d) &\equiv \mathbb{E}[y_{i2} | \alpha_{iA}, k_{i1} = A, D_A = 0] - \mathbb{E}[y_{i2} | \alpha_{iA}, k_{i1} = A, D_A = 1], \\ l_i^{\text{switch}} &\equiv l_i^{\text{switch}}(\alpha_{iA}, d) &\equiv \mathbb{E}[y_{i2} | \alpha_{iA}, k_{i1} = A, D_A = 0] - \mathbb{E}[y_{i2} | \alpha_{iA}, k_{i1} = A, k_{i2} = B, D_A = 1], \\ l_i^{\text{stay}} &\equiv l_i^{\text{stay}}(\alpha_{iA}, d) &\equiv \mathbb{E}[y_{i2} | \alpha_{iA}, k_{i1} = A, D_A = 0] - \mathbb{E}[y_{i2} | \alpha_{iA}, k_{i1} = A, k_{i2} = A, D_A = 1], \\ p_i &\equiv p_i(\alpha_{iA}, d) &\equiv \mathbb{P}(k_{i2} = B | k_{i1} = A, \alpha_{iA}, D_A = 1). \end{aligned}$$

We start with the simplest case, where occupational prices π_{kt} are revealed at the start of each period and there are no switching costs. Hence, occupational choice is a sequence of static decisions that can be analyzed in isolation. Panel (a) of Figure 1 illustrates occupational choices in the two periods as a function of workers' skills. The set of workers who start out in occupation A but then switch to B is indicated by the blue area in the figure. Given uniformly distributed skills, the figure shows that $\frac{\partial p_i}{\partial \alpha_{iA}} \leq 0$. We show in the online appendix that also $\frac{\partial l_i^{\text{switch}}}{\partial \alpha_{iA}} \geq 0$, and that $\frac{\partial l_i}{\partial \alpha_{iA}} > 0$: mean losses from occupational decline increase with initial earnings.

To understand the intuition for these results, call occupation A “typist” and occupation B “cashier”, where typists suffer a negative demand shock. The worst typists could only become the worst cashiers, otherwise they would have chosen to be cashiers in period 1. But the best typists can at most become the best cashiers, and in general they will not all be the best cashiers. Therefore, the best typists are less able to mitigate their earnings losses by becoming cashiers, and they suffer larger losses than the worst typists. This argument suggests that switching probabilities are decreasing and earnings losses are increasing in ability under a large set of alternative assumptions on the skill distributions.

Next, we assume there is a constant switching cost $c \in (0, d)$ for moving from occupation A to B . Occupational choice is no longer a period-by-period decision. Instead, workers choose in period 1 the occupation with the highest expected present discounted value of log earnings, net of switching costs. Let us assume that occupational log prices follow a random walk, $\mathbb{E}[\tilde{\pi}_2] = \tilde{\pi}_1 = 0$, where the last equality is due to our earlier simplifying assumption.¹⁰ Panel (b) of Figure 1 shows that occupational choices are qualitatively similar to the baseline model, except that the blue region marking the workers who switch is smaller than in panel (a). Again we have $\frac{\partial l_i}{\partial \alpha_{iA}} > 0$.

Instead of a constant switching cost, let us now assume that the cost for moving from A to B equals $C - \alpha_{iB}$, with $C > \bar{\alpha}$. This structure of switching costs captures in a reduced form way the frictions that occupational moves may entail: for example, job search may take time, and those more able in the new occupation may find a job more quickly. We continue to assume

¹⁰Instead of the random walk assumption we could impose that demand changes are somehow otherwise perfectly unforeseen, for instance due to adaptive expectations (in the online appendix we consider the case where demand changes are anticipated).

that occupational log prices follow a random walk. Panel (c) of Figure 1 shows that low-ability workers do not leave occupation A , and among high-ability workers, $\frac{\partial p_i}{\partial \alpha_{iA}} > 0$. We show in the online appendix that $\frac{\partial l_i^{\text{switch}}}{\partial \alpha_{iA}} < 0$ (taking into account earnings losses due to the time cost of switching), so that $\frac{\partial l_i}{\partial \alpha_{iA}} \leq 0$: mean losses from occupational decline (weakly) decrease with initial earnings.

In terms of the example above, in this case the worst typists do not switch, because their initial choice of occupation A reveals not only low earnings potential in occupation B but also a large switching cost. Among the best typists, however, many possess substantial earnings potential as cashiers, as well as low switching costs. Therefore, the best typists are on average better able to mitigate their earnings losses by becoming cashiers, and hence the earnings losses from the demand shock are smaller for the best typists than for the worst typists.¹¹

Building on the previous case, we now explore a version of the model that includes involuntary job displacement. Suppose that workers who start in occupation A experience job displacement with some probability that is independent of skill, and also incur a time cost $C - \alpha_{iA}$ to find a job in occupation A . Here we have in mind exogenous job losses, for instance due to plant closure, which are a standard feature of search models (see for instance Pissarides, 2000). There is a fraction of high-ability workers who switch occupation regardless of displacement. In addition, now a fraction of low-ability workers also switch, but only if they are displaced. This is illustrated by the yellow area in panel (c') of Figure 1. Moreover, the earnings losses experienced by these displaced movers are larger than those of comparable stayers. This is by revealed preference: a worker in the yellow region prefers to remain if not displaced, so her non-displaced counterpart (with the same period-1 earnings) necessarily incurs a lower earnings loss. We show in the online appendix that $\frac{\partial l_i}{\partial \alpha_{iA}} \leq 0$, as before. Unlike in the case without displacement, however, p_i is U-shaped in initial earnings. This is because the probability of a displacement-induced switch is decreasing, and that of a voluntary one is increasing in initial earnings. The earnings loss l_i is again decreasing in initial earnings, as the costs of moving

¹¹While our model excludes occupation-specific human capital, it does allow us to think about some of its potential implications. For example, if all workers accumulate occupation-specific human capital additively (in logarithms) the effects are similar to adding constant switching costs, since switching means foregoing this capital. And if higher ability workers accumulate more occupation-specific human capital they become less occupationally mobile, in contrast to the case of heterogeneous switching costs discussed above.

jobs—both within and across occupations—decrease with initial earnings.¹²

As a final variation on our model, we consider a case where period-2 prices are revealed to be $\tilde{\pi}_2 = d$ at the start of period 1. In the presence of switching costs, some workers that would otherwise have chosen occupation *A* in period 1 instead start out in occupation *B*. This means that the fraction of workers switching after period 1 is smaller, and it could even be zero if switching costs are large. Since there is less switching, earnings losses are larger than in the case of unanticipated shocks, for a given d .

We conclude this section by summarizing the main results from our model. The baseline frictionless model makes three predictions: the probability of leaving a declining occupation is decreasing in initial earnings; earnings losses due to occupational decline are increasing in initial earnings; and earnings losses of those who leave a declining occupation are less than the losses of those who remain. Anticipating that these predictions are inconsistent with our empirical findings, we consider how occupational switching costs can reconcile our results. Introducing an occupational switching cost that is decreasing in the worker’s earnings in the destination occupation, leads to a positive relationship between switching probabilities and initial earnings, and a negative relationship between earnings losses and initial earnings. Allowing for displacement, together with a cost of switching jobs within an occupation, implies that switchers’ earnings losses may be larger than those of stayers. Moreover, displacement can cause switching probabilities to be U-shaped in initial earnings, whereby low-earning workers switch involuntarily if displaced, while high-earning workers switch voluntarily regardless of displacement.

3 Data

Our main empirical analysis uses individual-level longitudinal administrative data covering the entire population of Sweden for several decades, and various editions of the Occupational Outlook Handbook (OOH) published by the Bureau of Labor Statistics (BLS). Some of our analysis

¹²We have also analyzed displacement under constant switching costs, that is, when workers incur a time cost $\hat{c} > 0$ to find a new job in *A*, or a cost c to find a job in *B*. This case is illustrated by panel (b’) of Figure 1, and details are given in the online appendix.

also uses the National Longitudinal Survey of Youth, which covers a sample of US residents, as well as other sources. Here we discuss key elements of the data we use, and leave many of the details to the online appendix.

3.1 Data sources

Our primary sources for measuring occupational decline are the 1986-87 and the 2018-19 editions of the Occupational Outlook Handbook (Bureau of Labor Statistics, 1986, 2018). The OOH describes the nature of work, the number of jobs, and the projected employment growth for hundreds of occupations. For a subset of these occupations, more details are reported, including (among much else) data on whether technology is expected to affect—or has already affected—the occupation in question, and if so in what way. In the 1986-87 edition, 401 occupations are described, covering about 80 percent of US employment. Detailed information is available for 196 of these occupations, covering about 60 percent of employment.¹³

Our main outcomes of interest come from Swedish micro data. We obtain basic demographic (year of birth, gender, education, and county of residence) and labor market (employment status, annual earnings, and industry) variables from the Integrated Database for Labour Market Research (LISA), a collection of administrative registers. For 1985-2013, LISA contains one observation per year for every individual aged 16-64 living in Sweden. Key variables, such as employment status and industry (as well as county of residence) are measured each November. We also use individual data from the Swedish Public Employment Service (PES), which contain information on the number of days registered as unemployed and number of days spent in retraining programs administered by the PES, for all individuals ever registered with the PES from 1992-2013.

To assess balance between treatment and control groups in terms of pre-determined characteristics, we use information on cognitive skills (an IQ-type measure) and non-cognitive skills (capturing psychological traits such as the ability to cope with stress) from the military enlistment. These data are described in detail by Lindqvist and Vestman (2011). We also use

¹³The number of distinct occupations in the OOH, as well as the number of occupations covered in detail, increased over time, so our crosswalk from the 1986-87 to the 2018-19 OOH is mostly, though not always, one-to-many.

information on parents' education and income from the 1985 version of LISA.

Our data on workers' occupations come from the population censuses, which were conducted every five years from 1960-1990, and from the Wage Structure Statistics (WSS) for the years 1996-2013.¹⁴ The WSS contains the population of public sector workers and a sample of about 50 percent of private sector workers. We apply sampling weights when working with the occupation variable from the WSS.

A useful feature of our data is that in the 1985 and 1990 censuses, workers' occupation is coded using a 5-digit classification, YRKE5, containing about 1,400 distinct occupations. This allows us to accurately merge occupation-level information from the US, as we describe below. Unfortunately, such detailed occupation codes are not available after 1990. From 1996-2013, a 3-digit classification containing 172 distinct codes, SSYK96, is available in the WSS. This classification is different from YRKE5, and the cross-walk between YRKE5 and SSYK96 likely introduces measurement error in workers' occupations after 1990. This limits our analysis of occupational employment shifts and individual workers' occupational mobility during 1985-2013.

Finally, we use information from the 1960 census, which allows us to calculate prior occupational employment changes at the 3-digit level using the YRKE3 classification, a coarser version of YRKE5 (there are 229 distinct codes that cover the period 1960-85).

While our main analysis focuses on outcomes in Sweden, we also use US panel data from the National Longitudinal Survey of Youth (NLSY 1979) to replicate the main analysis for the US. We leave the detailed discussion of these data to the online appendix.

3.2 Construction of key variables

To construct our measure of occupational decline we begin with the OOH data. Mapping occupations across the 1986-87 and 2018-19 editions of the OOH, we calculate the percentage growth in employment 1984-2016.¹⁵ If, after a careful search, a 1986-87 occupation has no counterpart in the 2018-19 edition, we classify it as having vanished, and assign a percent-

¹⁴We also use individual-level earnings data for 1975 and 1980 from the population censuses.

¹⁵The 1986-87 OOH reports employment for 1984, while the 2018-19 edition reports 2016 employment.

age growth of -100.¹⁶ While few occupations actually disappeared, some occupations declined sharply, including both white-collar (typists, drafters, and telephone operators), and blue-collar (precision assemblers, welders, and butchers) jobs.

We also record the projected employment growth of each occupation from the 1986-87 OOH. The BLS constructs these predictions using a careful and lengthy procedure.¹⁷ In the 1986-87 OOH, forecasts were reported in categories: “declining”, “little or no change”, “increasing slower than average”, “increasing about as fast as average”, and “increasing faster than average”. We create a cardinal predicted growth index assigning these categories the numbers 1-5 (where higher numbers correspond to more positive predicted employment changes). We report results both from using this index and using the categorical outlook variable.

In order to merge the OOH-based variables to the Swedish data, we map the 401 1986-87 OOH occupations to the 1,396 5-digit Swedish occupation codes available in the 1985 census. We successfully map 379 US occupations to 1,094 Swedish occupations—we are able to find corresponding US occupations for 91 percent of Swedish workers in 1985. We map percentage changes in US employment 1984-2016, as well as 1986-87 OOH predictions (categorical and index), to Swedish 5-digit occupations using our crosswalk, applying weights (OOH 1984 employment shares) in the case of many-to-one matches.

We define a Swedish 5-digit occupation as declining if the weighted employment growth of its corresponding OOH occupations is negative and larger (in absolute magnitude) than 25 percent. We regard this as a sensible threshold: smaller observed declines may result from measurement error from matching OOH occupations over time. At the same time, we report robustness checks using a number of alternative thresholds. We also use information from the OOH to determine whether technology likely played a role in the decline, as we further explain in the online appendix. In 1985, 13 percent of Swedish employees worked in subsequently declining occupations, and 8 percent worked in subsequently declining occupations where the decline is linked to technological change.

¹⁶Between the 1986-87 and 2018-19 editions of the OOH, some occupations were split or merged, which we take into account when computing the percentage growth. See the online appendix for details.

¹⁷Veneri (1997) evaluates the ex-post accuracy of the projections used in the 1986-87 OOH, and concludes that it correctly foresaw most occupational trends, although there were non-trivial cases of error.

We construct several left-hand side variables that characterize workers' career outcomes spanning the years 1986-2013, that is, starting with the first year after we measure treatment and ending with the last year available in our data. We start by summing up years observed as employed and real annual labor earnings, obtaining the variables cumulative years employed and cumulative earnings. Following Autor, Dorn, Hanson, and Song (2014), we measure normalized cumulative earnings, which is the ratio of cumulative earnings to predicted initial earnings.¹⁸ We consider further earnings measures—such as rank, discounted cumulative earnings, and earnings growth—in robustness checks.

Our measure of long-run occupational mobility is a dummy variable that equals one if the individual worked in the same 3-digit SSYK96 occupation in 2013 as 1985. It equals zero if the individual works in a different occupation or is not employed.¹⁹ Using the PES data, we calculate cumulative days spent unemployed and cumulative days spent in retraining during 1992-2013. We define dummy variables for ever unemployed and ever having participated in retraining. Finally, we calculate the retirement age, where we define retirement as a continuous spell of zero annual earnings up to and including age 64.²⁰

3.3 Sample restrictions

Our starting sample contains all individuals born between 1921-1969 and hence aged 16-64 (at some point) in 1985; who were employed in November 1985; whose annual earnings in 1985 were no less than the “base amount” (Swedish: *basbelopp*) specified by the social security administration; and for whom we have the relevant demographic and labor market information.²¹

¹⁸The prediction comes from a regression of log earnings on a quartic in age and dummies for gender, county, and seven education categories, run separately for each 3-digit SSYK96 occupation in 1985. We divide by predicted rather than actual initial earnings to eliminate transitory earnings variation, which would introduce an important role for mean reversion into the distribution of normalized cumulative earnings. Autor, Dorn, Hanson, and Song (2014) divide cumulative earnings by earnings averaged across four pre-treatment years for the same reason. Since we do not have annual earnings information prior to 1985, we normalize by predicted earnings instead.

¹⁹Our measure of occupational mobility does not capture any temporary exits during the intervening years if workers returned to their initial occupation. A limitation of our data is that they are not conducive to studying high-frequency occupational mobility: During the years 1986-1989 and 1991-1995, we do not observe workers' occupation. And during 1996-2004, the SSYK96 variable contains substantially fewer distinct codes than from 2005 onwards.

²⁰The LISA database includes individuals older than 64 only during later years. Since we do not consistently observe individuals beyond age 64, we assume for all years that individuals aged 65 or older have retired.

²¹The base amount is used as an accounting unit when calculating benefits, and it is typically equal to about three months' worth of full-time work at the median wage. As we do not observe hours worked or fulltime status,

There are 3,061,051 individuals fulfilling the above criteria.²² Our *baseline sample* further restricts the sample to those aged 25-36 in 1985. We drop younger workers, who are less likely to have settled on an occupation. And we drop middle-aged and older workers from our baseline sample because we want to focus on workers who did not reach retirement age by 2013, the end of our period of study, in our main analysis. We analyze these older workers separately.

4 Empirical strategy

4.1 The estimating equations and their interpretation

Our objective is to estimate the consequences of occupational decline for individual workers' careers. Consider a regression of cumulative career outcomes—such as cumulative years employed, or cumulative earnings—on an indicator for working in 1985 in occupations that subsequently declined, conditional on a set of controls. The probability limit of the regression coefficient on the declining indicator can be expressed as a difference in conditional means, which in turn can be decomposed into the sum of a treatment effect and selection bias,

$$\begin{aligned}
 & \underbrace{\mathbb{E}[y_{ik2}|k_{i1} \in A, D_A = 1, x_{i1}] - \mathbb{E}[y_{ik2}|k_{i1} \in B, D_B = 0, x_{i1}]}_{\text{Observed difference in conditional means}} = \\
 & \quad \underbrace{\mathbb{E}[y_{ik2}|k_{i1} \in A, D_A = 1, x_{i1}] - \mathbb{E}[y_{ik2}|k_{i1} \in A, D_A = 0, x_{i1}]}_{\text{Effect of occupational decline; } -\mathbb{E}[l_i] \text{ in model}} \quad (2) \\
 & \quad + \underbrace{\mathbb{E}[y_{ik2}|k_{i1} \in A, D_A = 0, x_{i1}] - \mathbb{E}[y_{ik2}|k_{i1} \in B, D_B = 0, x_{i1}]}_{\text{Selection bias}}.
 \end{aligned}$$

Here, y_{ik2} is the outcome of interest, such as cumulative earnings of individual i (who is employed in occupation k in 1985) in period 2 (1986 through 2013).²³ Our notation separates declining occupations (A) from non-declining ones (B). D is an indicator for occupational de-

we use the base amount to exclude individuals with little labor market attachment.

²²There were 5,281,382 individuals aged 16-64 in Sweden in 1985. Of those, 4,186,512 were employed in November 1985, and among them, 3,648,034 earned no less than the base amount during 1985. The reduction to 3,061,051 is due to missing education, industry, or occupation information, including cases where YRKE5 occupations do not have matches in the OOH.

²³In this paper we focus on labor market outcomes, and in companion work we study how occupational decline matters for other socio-economic outcomes, including health, family composition, geographic location, and welfare transfers.

cline, which allows us to consider the hypothetical situation where occupational decline did not take place. We motivate our use of an indicator for occupational decline in Section 4.2 below. x_{i1} is a set of (yet unspecified) controls, which we also revisit below.²⁴

In our model, the selection term equals zero due to the symmetry assumptions we make, and there is no need for controls. In practice, the identifying assumption for the regressions without controls is too strong, because the selection term reflects both differences in individuals sorting across occupations as well as differences between occupations A and B even in absence of occupation decline. Our empirical strategy aims to mitigate both of these types of selection concerns.

Our first step towards addressing the sorting of individuals is to control for a rich vector of individual characteristics in period 1, that is, in 1985: gender, age, educational attainment and county of residence dummies, and earnings. But a natural question is whether the sorting of workers makes those in declining occupations differ in unobserved characteristics which may affect earnings trajectories, such as their cognitive and non-cognitive ability, even conditional on these controls. Fortunately, we can (and do) address this concern using measures of both cognitive and non-cognitive skills from the military enlistment, which are available for men of a subset of cohorts for whom military service was compulsory. We also check whether workers in declining occupations differ in other background characteristics, namely parental education and pre-1985 earnings.

But even when the vector of controls x_{i1} ensures balance in terms of individuals' unobserved characteristics, the selection term will be non-zero if earnings growth in occupations A and B would have been different in the absence of occupational decline—that is, if a worker's occupation affects her earnings growth even without occupational decline. To mitigate this concern, we use data from 1985 to estimate expected occupation-specific lifetime earnings profiles, which we add as controls to the regressions.

To further control for differences between declining and non-declining occupations, we use

²⁴As in our model, the effects that we estimate are on the workers in declining occupations relative to other workers, since we cannot identify the aggregate effects of occupational decline. We do, however, explore the role of one mechanism through which occupational decline may affect non-declining occupations, namely the movement of workers into non-declining occupations, as we further discuss in the next section.

information from the 1986-87 OOH. The BLS authors went to great lengths to accurately forecast occupational employment changes. Once we condition on predicted occupational growth, we likely remove much of the remaining differences between declining and non-declining occupations. Since the occupational decline and the forecasts that we use rely on US data, we also control for each occupation's level of employment and pre-existing employment growth trends in Sweden. In some specifications we use two additional sets of controls: broad (1-digit) occupation dummies and (2-digit) industry dummies. Adding these controls comes at the cost of reducing the variation in occupational decline, since it only uses variation in occupational decline between very similar occupations. Nevertheless, by gradually adding controls we achieve plausible bounds on the estimated effects of occupational decline, as we further discuss in the next section.

Taken together, our estimating equation takes the form

$$y_{i2} = \beta D_{k_{i1}} + \gamma x_{i1} + \delta w_{k_{i1}} + \varepsilon_i, \quad (3)$$

where $D_{k_{i1}}$ is an indicator for working in 1985 an occupation that subsequently declined; x_{i1} is a vector of individual characteristics, measured in 1985, as discussed above; w_k is a vector of occupational characteristics; and ε_i is the error term, which we conservatively cluster by three-digit Swedish occupations.²⁵

We provide further evidence that our identification strategy plausibly addresses the selection issues by considering two additional sets of outcomes. First, we examine the earnings of workers in the years before the occupational decline that we study. Second, we study the cumulative earnings during the first few years our study, when the effect of occupational decline were likely limited.

A different question regarding our approach is whether occupational decline that is specifically linked to labor-replacing technologies has distinct consequences from demand-driven occupational decline in general. To provide evidence on the role of technology, we restrict some of the analysis to occupational declines that are explicitly linked to concrete new technologies,

²⁵As discussed above, some specifications also control for industry fixed effects.

such as personal computers and robots.

Yet another question about the interpretation of our estimates is whether they can be generalized beyond Sweden. To shed light on this question, we repeat our main analysis using data on US workers (NLSY 1979), which allows us to study the consequences of occupational decline in the US context, albeit with more data limitations.

One final step in our empirical analysis is to consider how the costs of occupational decline may fall differently on workers with different initial earnings rank within their occupations, in line with the discussion in the model section.

4.2 Rationale for measuring occupational decline using US data

Prior literature has documented that shifts in occupational employment are strongly correlated across countries, see for instance Goos, Manning, and Salomons (2014) documenting job polarization across European countries, and in particular Adermon and Gustavsson (2015) on job polarization in Sweden. Here we explain why using measures of occupational decline from the OOH is not only feasible, but also desirable.

We begin by explaining why we prefer this measure of decline to an alternative measure using the SSYK96 codes. First, there are 401 OOH codes compared to just 172 (three-digit) SSYK96 codes, and having more codes affords us more variation from small and declining occupations. For example, it lets us separate typists, whose employment fell sharply, from secretaries, whose employment grew. To use the OOH codes we match them to YRKE5 codes, but since the YRKE5 are more numerous we do not lose much variation. Second, since the SSYK96 codes were introduced from 1996 they reflect a judgement on an occupation's importance made after the start of the occupational decline that we study. Consequently, SSYK96 are more likely to pool occupations with low employment in 1996 (including declining ones) with non-declining occupations. Because the 2018-19 OOH separately describes even occupations with very low employment, this is less of a problem for our approach. Finally, using occupational declines measured in Sweden as a regressor where the dependent variable is change in earnings creates a problem of simultaneity. This problem is mitigated by using the OOH

measure.

At this stage readers may also ask: why do we report reduced form results using the OOH decline measure rather than use it as an instrument for occupational decline measured in Sweden using SSYK96? Our rationale for the reduced form approach is that it preserves much more of the variation that we are interested in, for several reasons. First, as noted above, if we use measures based on SSYK96 codes, we lose much of the variation in occupational decline because of the coarseness of the classifications and the lower likelihood of separating occupations in decline. Second, 2SLS would exacerbate this problem, since it only uses part of the variation in the decline. Finally, as we show below, while we still have power to detect changes in occupational decline in Sweden, once we control for predicted changes we are left with a weak instrument.

Still another question is why we focus only on occupational declines instead of using the full variation in OOH occupational change. Again there are several factors that influence our choice. First, declines are interesting from the perspective of their social costs and policy implications. Second, large declines in employment are likely driven by declines in labor demand, whereas increases in employment may also reflect shifts in labor supply. Finally, as we explain below, we use different cutoffs in the regressions as well as graphical evidence to show that the costs of occupational change are concentrated among those who experience substantial occupational declines; increases or moderate declines seem to matter little relative to each other. Nevertheless, for completeness we also report estimates using the full variation in occupational changes.

To conclude, we note that while our reduced form estimates on their own do not deliver immediately interpretable magnitudes, we are able to assess the quantitative importance of say, estimated earnings losses, by relating them to the estimated impacts on occupational mobility, and also, to the difference in employment growth between declining and non-declining occupations. Of course, our discussion above suggests that our estimates on occupational change and mobility in Sweden may understate the true extent of these changes, since they rely on the SSYK96 classification.

5 Empirical analysis

In this section we present the findings from our empirical analysis. First, we quantify occupational decline in Sweden and discuss sorting into declining occupations. Second, we study how employment, earnings, and occupational mobility differed for workers in declining occupations. Third, we investigate how the consequences of occupational decline differed by workers' initial within-occupation earnings rank. Fourth, we explore some of the mechanisms through which occupational decline operates, including unemployment, retraining, and early retirement. Fifth, we examine whether occupational declines with observed links to technology has distinct consequences. Sixth, we repeat the main elements of our analysis using NLSY panel data from the US. At the end of this section, we interpret our findings through the lens of the theoretical model from Section 2.

5.1 Occupational decline and sorting across occupations

We begin by quantifying workers' exposure to occupational decline. In Table 1 we report estimates of equation (3), where the dependent variable is log employment change from 1985-2013 in each worker's three-digit occupation. Panel A shows estimates for workers aged 16-64 in 1985, and Panel B focuses on our main sample of workers—those aged 25-36 in 1985. The results, which are similar across panels, suggest that workers in declining occupations are exposed to a log employment change that is about 50 log points lower than for workers in non-declining ones; about 40 log points lower when we compare observationally similar workers; and lower by about 20-30 log points when we also include occupation and industry controls. It is important to keep these results in mind when interpreting our findings from putting individual-level career outcomes on the left-hand side of equation (3). While adding more controls reduces the risk of omitted variables bias, the results in Table 1 show that this also leaves less variation in exposure to occupational change. We also note, as discussed in Section 4.2, that these estimates likely understate the employment decline for 5-digit occupations, which we are unable to

measure.²⁶

Table OA1 reports similar estimates, aggregated by three-digit occupations and weighted by 1985 Swedish employment shares, using our main sample of workers. This table again shows large declines in Swedish employment in occupations we classified as declining based on the OOH data, in particular, a difference of about 75 log points between all-declining and none-declining 3-digit occupations (column (1)).²⁷ Table OA1 also explores the extent to which declines were predictable in 1985. It shows that the OOH-based predicted growth index has strong explanatory power for Swedish employment growth (column (3)). Entering the OOH predictions as categorical variables only marginally improves the forecast, as seen in column (4). Finally, the difference in employment growth between all-declining and none-declining 3-digit occupations is still about 45 log points when controlling for initial Swedish employment shares, prior Swedish employment growth, and the OOH predictions, as seen in columns (5)-(6).

Having described the extent of occupational declines, we turn to the sorting of individuals in 1985 into subsequently declining occupations. Table 2 presents results from regressions of several individual characteristics on an intercept and the declining indicator. The top panel considers the entire working-age population with non-missing demographic and labor market information, and the bottom panel focuses on our main sample. In both cases, the sorting patterns are similar: those in occupations that subsequently declined were of similar age, and more likely to be male, less educated, and more likely to be employed in manufacturing. Coincidentally, the gender gap in earnings is offset by the differences in schooling, and on net, the workers in subsequently declining occupations had similar earnings to others in 1985.

We next investigate whether there is sorting into declining occupations based on cognitive skills, non-cognitive skills, and parental attributes, and if so, whether any differences in these

²⁶The difference of 50 log points translates into an employment decline of about 18 percent. Let y_i be the log employment change assigned to each individual based on her 1985 5-digit occupation, and D_i be the declining indicator. From the regression $y_i = \alpha + \beta D_i + \varepsilon_i$ we obtain $\mathbb{E}[\exp\{y_i\} | D_i = 1] = \exp\{\alpha + \beta\} \mathbb{E}[\exp\{\varepsilon_i\} | D_i = 1]$, and plugging in our estimates, $\exp\{0.06 - 0.49\} \times 1.26 = 0.82$. As we argue in the text, this likely understates the actual average decline in the 5-digit occupations.

²⁷The difference in the coefficients between the first columns in Tables 1 and OA1 is to be expected, because the micro data equivalent of a regression using aggregate data in cases like this is a two-stage least squares (2SLS) regression, not an OLS regression. If in the micro data we instrument the declining indicator with the complete set of 3-digit occupation (SSYK96) dummies, we obtain a coefficient identical to that in column (1) of Table OA1. However, the results from the OLS regression are easier to interpret and to relate to our results from putting individual-level outcomes on the left-hand side.

variables disappear once we control for the individual characteristics described in the previous paragraph. Columns (1) and (4) in panel A of Table 3 show that in 1985, the cognitive and non-cognitive skills of men in subsequently declining occupations were lower than those of the other men by about 0.2-0.25 standard deviations. However, these differences shrink and become insignificant once we add the individual controls (columns (2) and (5)), and are essentially eliminated when we also include industry and occupation controls (columns (3) and (6)). Panels B and C repeat the analysis for mothers' and fathers' schooling and earnings in 1985, which have the advantage of being available for women as well as men. The pattern is qualitatively very similar to that of the military skill measures: working in a subsequently declining occupation is associated with lower parental schooling and earnings, but these differences disappear once adding controls. In sum, Table 3 suggests that, although there is negative selection into occupations that later decline, most of this selection can be mitigated using suitable controls such as education and 1985 earnings. To the extent that minor negative selection persists in the regressions with individual controls, these regressions may slightly overstate the negative effect of occupational decline on workers, but this is not a concern once we add all the controls.

As a final check for sorting into declining occupations, we investigate earnings in 1980 for the older cohorts in our baseline sample (individuals aged 25-31 in 1980).²⁸ Again we find that conditional on individual-level characteristics, there are essentially no differences in prior earnings, as seen in Figure OA2. Taken together, the results in this section suggest that concerns about sorting into declining occupations are largely alleviated when we include suitable controls.

5.2 Main results on employment, earnings, and occupational mobility

Table 4 reports results from estimating equation (3) using our main sample of workers aged 25-36 in 1985. Panel A shows that workers in declining occupations spent about nine fewer months (0.73 fewer years) in employment from 1986-2013 (column (1)). Once we add individual controls, this estimate reduces to about six months, or about two percent of the sample mean of

²⁸Earnings data for the population of Swedish workers are not available at annual frequency prior to 1985. We obtain prior earnings from the population censuses, which were carried out every five years until 1990.

about 23 years (column (2)). Next, we add more controls and compare those who experienced occupational declines to observationally similar workers in similar occupations and industries. These specifications suggests that the losses from occupational decline averaged about two months (0.2 years) of employment, or about one percent of the sample mean (columns (3)-(6)).

Panel B of Table 4 reports results from using cumulative earnings 1986-2013 as the outcome. Column (1) shows that working in a declining occupation was associated with 350,000 Swedish Krona (SEK) lower cumulative earnings, or about 5 percent of the sample mean.²⁹ When including individual controls, the estimated loss is similar, though the confidence interval is much tighter (column (2)). Further adding occupational controls cuts the loss to less than two percent of the sample mean.

In panel C we examine earnings losses from occupational decline using an alternative earnings measure: cumulative earnings divided by predicted initial earnings (see Section 3.2 for details on the construction of this variable). Depending on the controls included, the estimated losses in cumulative earnings range from around 100 to 220 percent of initial annual earnings, or from 2.5-5.7 percent in terms of the sample mean, quite similar to the results in panel B.³⁰

In Figure 2, we present a dynamic counterpart to the results reported in panel B, columns (2) and (6) of Table 4. Here we use as outcomes each year's earnings and cumulative earnings from 1986 up to the year indicated on the horizontal axis of each chart. The top right panel of Figure 2 is suggestive of a smooth process of occupational decline, with earnings losses building up gradually. However, the top left panel reveals that losses in annual earnings suffered by workers in declining occupations were larger during the 1990s and late-2000s recessions. The picture is similar when we divide the coefficients by the mean of cumulative earnings at each horizon (bottom panels). As before, the losses are smaller when we include occupation and industry controls.

Next, we investigate occupational mobility. Table 5 reports estimates of equation (3) with

²⁹We inflate all SEK figures to 2014 levels. Average annual earnings of Swedish workers, conditional on being employed in November and earning at least the base amount during the year, were SEK190,200 in 1985 and SEK330,800 in 2013, in terms of 2014SEK. We do not express these amounts in USD due to exchange rate fluctuations. For instance, SEK1,000 were worth about USD150 in January 2014, but about USD130 in December 2014, and about USD110 in October 2018.

³⁰Below we discuss results using alternative functional forms for cumulative earnings.

indicators for working in 2013 in the same occupation as in 1985 (or in a similar ones) as outcomes. As we do not want to condition the sample on being employed in 2013 (which is also an outcome), we have that ‘not remaining’ in the same occupation could reflect either occupational switching or non-employment, a point to which we return below.

Column (1) in panel A of Table 5 shows that the probability of remaining in the same 3-digit occupation was around 14 percentage points lower in declining occupations, compared to a mean of 29 percent in our sample. In other words, by 2013 a little over 70 percent of all workers had left their 1985 occupations (or left employment altogether), and the probability of staying in the same occupation was roughly halved for those starting in declining occupations. When we compare observationally similar workers, occupational decline appears to reduce the probability of remaining in the 1985 occupation by 11 percentage points, and when further restricting the comparison to similar occupations and industries, the estimate falls to 4.5 percentage points. Panels (B) and (C) of Table 5 show similar, albeit somewhat smaller, coefficients when we look at the probability of remaining in more broadly defined (2-digit or 1-digit) occupations. It is noteworthy that even when we consider 1-digit occupations, only about 40 percent of the sample remained in the same broadly defined occupation over the 28-year period that we study.³¹

Having presented our main results on career employment, career earnings, and occupational mobility, we now return to the issues of interpretation alluded to in Section 4.1. Our results from including individual, occupation, and industry controls (column (6) in Tables 4 and 5) plausibly provide us with conservative estimates of the losses from occupational decline—about 2 percent of mean cumulative earnings over 28 years—given the balance of pre-determined characteristics conditional on these controls, as well as the fact that no earnings losses appear in the first 5-10 years (Figure 2). As argued in Section 5.1, the specification only controlling for individual characteristics (column (2) in Tables 4 and 5) may slightly overstate the losses from occupational decline—5 percent of mean cumulative earnings—as it leaves minor differences in some of the pre-determined characteristics. In addition, Figure 2 shows earnings losses based on this specification already in the years immediately after 1985. On the other hand, we have

³¹For related discussions of the importance of switching occupations in the presence of technological change, see Cortes (2016) and Caselli and Manning (2018).

also seen that the variation in occupational decline is much reduced when including occupation and industry controls (columns (2) and (6) in Table 1). While larger declines going together with larger losses and mobility is consistent with our theoretical model, we note that the reduction in employment declines is about 50 percent between the two specifications in Table 1, whereas the coefficients for individual employment, earnings, and mobility mostly decline by more than 50 percent. In conclusion, we consider 2-5 percent of mean cumulative earnings to be a credible range of estimates of the losses from occupational decline.

5.2.1 Robustness of main results

Our first set of robustness checks relates to the choice of functional form of occupational decline. The declining indicator is based on a 25-percent cutoff, conservatively identifying occupations whose (US) employment fell substantially since the mid 1980s. We also explore a range of alternative cutoffs and find that higher cutoffs (in the sense of isolating larger employment declines) usually result in larger estimated losses and mobility responses. In addition, our results are very similar when we exclude occupations that grew rapidly from the control group (see Table OA2 for both sets of results). While we focus on a binary definition of occupational decline as motivated in Section 4.2, we also explore the relationships between our key outcomes of interest and the full variation in US and Swedish employment growth. The (residualized) associations of cumulative earnings and occupational mobility with occupational employment growth are mostly flat apart from a drop in occupations that declined substantially (see Figures OA3 and OA4, and corresponding regression results in Table OA3).

With a second set of robustness checks we aim to address the risk of underestimating the losses from occupational decline. The magnitude of our estimated earnings losses may be understated because workers who leave declining occupations flock to similar non-declining occupations, depressing the wage in these ‘control’ occupations. And even in the absence of such general equilibrium effects, employing a rich set of controls may cause us to put more weight on groups of comparable occupations where there are roughly as many declining as non-declining sub-occupations. In such cases, workers may have many substitute occupations to choose from,

so our estimates may understate the true average treatment effect. To mitigate these concerns, we run what we refer to as ‘doughnut’ specifications, namely the same regressions as those we report in Tables 4 and 5 but excluding 3-digit (SSYK96) occupations in which some but not all 5-digit occupations are declining. We indeed estimate larger earnings losses than in our baseline specifications, but only slightly so, ranging from 3-6 percent of mean earnings (see Tables OA4 and OA5).

Finally, our conclusions about earnings losses are robust to using different functional forms of earnings, as we discuss further in the next sub-section.

5.3 Heterogeneity by within-occupation earnings rank

We now examine how employment and earnings losses from occupational decline varied by initial within-occupation earnings rank. We estimate equation (3) allowing the coefficient on the declining indicator to vary by earnings rank, and report the results in Table 6. Panel A shows that lower ranked workers suffered larger employment and earnings losses than average as a result of occupational decline (columns (1)-(6)): the coefficients on the interaction of the declining indicator with earnings rank are positive and precisely estimated. Moreover, these estimates barely change when we add occupation and industry controls over individual-level controls, though the main coefficients on the declining dummy—giving the employment and earnings loss for the median worker—are affected by the inclusion of additional controls. The magnitudes implied by the interaction coefficients are meaningful and imply, for instance, that compared to the 25th-percentile, the 75th-percentile worker suffered a 5-percent lower employment loss and a 6.5-percent lower earnings loss (both in terms of the overall mean).

This pattern is robust to alternative specifications that replace the linear rank measures with dummies for the bottom and top terciles. This specification also allows us to characterize losses for low-ranked workers directly. Panel B of Table 6 shows that workers at the bottom tercile of their starting occupations’ earnings distributions suffered employment losses of 1.2-1.4 years (5.5-6.5 percent of mean employment in the bottom tercile) and earnings losses of around 8-11 percent of bottom-tercile mean earnings. Indeed, the estimates of mean losses reported in the

previous sub-section mask more substantial losses for low earners (within an occupation). Our findings about earnings losses are robust to a number of alternative ways of measuring career earnings, as shown in Table OA6.³²

The pattern for the probability of remaining in the initial occupation appears to be non-monotonic: among the workers in declining occupations, both bottom-tercile and top-tercile workers were less likely to remain in their starting occupations (panel B, columns (7)-(8)). These interaction coefficients are larger than ten percent of the overall mean (although in the case of the top tercile, not precisely estimated). This hump-shaped pattern of staying probabilities (U-shaped in exiting probabilities) is intriguing from a theoretical point of view, as we discuss below.

One potential challenge in interpreting the results of Table 6 is that those with low earnings in their occupation may have differed from others along some observable dimensions, such as gender, age, or geography. To mitigate this concern, we re-estimate the regressions using workers' within-occupation rank in residualized earnings, where the residuals come from a regression of earnings on female, cohort, and county-of-residence dummies. As Table OA7 shows, in terms of employment and earnings losses the results are qualitatively unchanged, and the magnitude of the interaction coefficients is only slightly reduced. However, using the residual-based rank measure, there is less support for the conclusion that bottom-ranked workers were less likely to remain in the initial occupation.

We now briefly examine earnings losses separately for those who remained in their initial occupation and those who did not. This purely descriptive exercise is motivated by the prediction of our baseline model in Section 2 that leavers should have lower losses than stayers. We estimate equation (3) with cumulative earnings as the outcome variable, and add on the right-hand side a dummy for having remained in the initial occupation, as well as its interaction with the declining dummy. Panel A of Table OA8 shows that among all workers, those who remained in their initial occupation had higher cumulative earnings, though in panel B we

³²We consider discounted cumulative earnings, applying a 5-percent discount rate; discounted cumulative earnings normalized by initial earnings; the percentile rank in cumulative earnings; the log of cumulative earnings; and the percentage change in earnings 1985-2013. As expected, the estimated losses in terms of discounted cumulative earnings are somewhat smaller at 1.5-4.5 percent of the overall mean, depending on controls, as more weight is put on earlier years in the career.

restrict the sample to those who were employed in 2013, and the finding reverses.³³ Importantly, in neither case is there evidence that those who remained in declining occupations did significantly worse than those who left a declining occupation. The same result holds when we focus on the bottom third (in terms of within-occupation earnings), see panel C. We discuss the interpretation of these results in light of the model in Section 5.7 below.

5.4 Unemployment, retraining, and early retirement

A natural question at this stage is to what extent the loss in years of employment due to occupational decline is accounted for by increased unemployment and retraining; as discussed above, data on these last two outcomes are available for the final 22 years of our study. Table 7 reports estimates using the main specifications from Tables 4 and 6 but this time using cumulative days of unemployment (panel A) and state-sponsored retraining (panel B) as outcome variables. Columns (1)-(4) of Panel A show that workers who started out in later declining occupations were only very slightly more likely to ever be unemployed, and columns (5)-(8) suggest that these workers accumulated 20-50 more unemployment days, though the estimates with more controls are imprecise. However, we again find substantial heterogeneity, with bottom-tercile workers in declining occupations spending 63 days more in unemployment, a substantial 20 percent of the mean.

Columns (1)-(4) of panel B suggest that occupational decline increased the risk of ever enrolling in state-sponsored retraining by 9-27 percent. The estimates for cumulative days spent retraining are similarly substantial, at least in relative terms (columns (5)-(8)). Our most conservative specification including all controls suggest that the median worker spent six more days in retraining, which amounts to 21 percent of the mean (ten days and 29 percent for the bottom-tercile worker).

Our estimates for unemployment and retraining can only explain part of the estimated employment losses. For bottom-tercile workers, we conservatively estimate an employment loss

³³Workers classified as having remained are employed in 2013 by construction, whereas those classified as not having remained might not have been employed in 2013 and thus have zero earnings in that year, and possibly in preceding years also.

of 1.16 years.³⁴ Of these, unemployment and retraining account for only 22 percent.³⁵ The remaining employment loss may be accounted for by job search that is not covered by unemployment benefits; private retraining; or time spent outside the labor force. Unfortunately, we lack the data to investigate this further.

There is however a group of workers for whom we are able to investigate the relationship between occupational decline and exit from the labor force, namely, older workers. Recall that workers in our baseline sample reached a maximum age of 64 in 2013. We now examine employment, earnings, and retirement for two groups of older workers, most of whom reached the usual retirement age of 65 well before the end of our sample period.

Panel A of Table 8 considers workers who were aged 37-48 in 1985. The employment losses among this group are a little larger than for our baseline sample: about 8 months (4 months) of a year of employment in the specification with individual (all) controls, or just under 4 percent (2 percent) of the group mean. About half of these employment losses are accounted for by a slightly younger age of retirement for those in declining occupations. The estimated earnings losses from occupational decline—about 6 percent (1.5 percent) with individual (all) controls—are similar to those of the baseline group. Finally, for this group we also find positive and significant interactions of the declining dummy with initial occupational earnings rank, suggesting once more that those who earned least within their occupation to begin with lost more years of employment from occupational decline.

Panel B of Table 8 suggests that for an even older group, those aged 49-60 in 1985, the occupational decline that we measure had more modest costs compared to the baseline group. This likely reflects the fact that we are measuring occupational decline over a longer period, and that these older workers had little exposure to the decline.³⁶

³⁴From panel B, column (2) in Table 6 we obtain $-0.03 - 1.13 = -1.16$. To complete the calculation, we divide the unemployment and retraining coefficients by 365 to get years, multiply them by 28/22 to account for the fact that these variables are only available during 1992-2013, sum them, and divide by 1.16.

³⁵Of the mean employment loss, unemployment and retraining explain about a third and a tenth of the time respectively.

³⁶We verify that for the groups of middle-aged and older workers, our declining indicator does not predict differences in prior earnings (1975 and 1980) conditional on controls, see Figure OA2.

5.5 Technology-related occupational decline

Consistent with much of the literature (Goos, Manning, and Salomons, 2014) we expect technological change to be a key driver of occupational decline, and especially occupational decline that is common to the US and Sweden. Nevertheless, there could be other drivers, including changes on the supply side (changes in demographics, trade shocks, or changes in government policy) and in consumer demand. Bearing this in mind, we now focus on occupations that are likely to have declined due to the introduction of labor-replacing technology, based on information from the OOH, as described in Section 3.2.³⁷

We find that workers' exposure to declines in Swedish occupational employment is of very similar magnitude regardless of whether we consider all occupations classified as declining, or only the ones we linked to technology (Table OA10, panel A, and Table OA11).³⁸ Moreover, technology-related occupational declines are not significantly different from other occupational declines in their implications for years of employment, cumulative earnings, and the probability of remaining in the initial occupation. One way to see this is by adding an indicator for technology-linked decline to equation (3). We find that the coefficients on this indicator are statistically indistinguishable from zero (columns (1) and (2) in panels B-D of Table OA10). Alternatively, considering technology-related declines on their own, we see very similar point estimates, both for the main effect and for the interaction with earnings rank, as for the full set of declines (columns (3)-(5)).

5.6 Studying occupational decline in the US using NLSY data

An important question is to what extent the magnitudes of the earnings and employment losses that we estimate are specific to Sweden and its institutional setting. We aim to answer this question using US data from the National Longitudinal Survey of Youth (NLSY 1979). In this analysis we try to stay as close as possible to the specifications we estimate for Sweden, but

³⁷Some of what we classify as technology-related decline may still be influenced by other factors, and we cannot rule out that technology played a role in the remaining declining occupations.

³⁸Workers starting out in 1985 in subsequently declining occupations, where we were able to identify a link to technology, were statistically indistinguishable from those in the remaining declining occupations, as seen in Table OA9.

some changes are necessary due to data limitations. The NLSY cohorts are younger than those we study in Sweden, so we set 1987 (instead of 1985) as the base year. This way it is still reasonable to use the same OOH data that we use for Sweden while allowing the youngest workers to have reached age 22 in the base year. This means that the cohorts we study in the NLSY are likely less attached to the labor force, but for the most part are likely to have completed college (if taking any). The geographic information in the NLSY is also limited, so we use region dummies instead of county dummies as controls. To ensure a sufficient sample size, we use the 1980 US census to construct occupational life cycle earnings profiles, and where necessary we impute earnings for years where they are not reported. Other aspects of the NLSY are discussed in the online appendix.

To shed light on how occupational decline shaped earnings in the US, Table OA12 reports estimates of specifications similar to those in panel B of Table 4. We estimate different specifications, and in all cases our point estimates are close to zero. While the estimates are quite imprecise, in our preferred specification the 95-percent confidence intervals exclude losses of 7 percent or more. We note that this is larger in magnitude than our main point estimate for Sweden, but a little smaller than the point estimate for the bottom tercile in Sweden.

There are several possible reasons why the NLSY estimates may be less precise than those we obtain using the Swedish data. First, the NLSY sample is only a small fraction of the US population, and the sample size is around two orders of magnitude smaller than the Swedish data. Second, workers in the NLSY were on average younger in the base year, and therefore may have been less attached to their starting occupation. Third, NLSY earnings are self-reported, while those in Sweden come from administrative records. Fourth, the NLSY suffers from more attrition and non-reporting compared to Sweden's administrative data. Finally, there may be other aspects of measurement that differ across the two countries (such as the measurement of occupations).^{39,40}

The picture is similar for cumulative weeks employed and unemployed as outcomes. The

³⁹We also do not find significant interactions with occupation-specific rank, likely because of a lack of power given the much smaller NLSY sample.

⁴⁰Figure OA5 uses the US NLSY data to repeat (as closely as possible) Figure 2 for Sweden. In the US, like in Sweden, the confidence intervals widen a little over time, but the US data show no clear trend for the point estimates.

estimates are for the most part imprecise, and the point estimates suggest that those in declining occupations spent a slightly larger fraction of their time in both employment and unemployment (compared to non-employment). The 95-percent confidence intervals exclude losses of 2 percent or more, which is again broadly consistent with our findings for Sweden (see Table OA13). Finally, we find some suggestive evidence that occupational decline made it less likely that workers remained in their 3-digit starting occupation (Table OA14).

5.7 Interpreting our findings through the lens of the theoretical model

We now discuss how our results relate to the insights from the theoretical model presented in Section 2. Our model assumes that occupational decline results from adverse demand shocks, so that affected workers suffer relative earnings losses and are more likely to exit their occupations. In our empirical analysis, we confirm that occupational decline was indeed associated with earnings losses and higher exit rates. Our results therefore support our interpretation that the occupational decline that we study was largely driven by changes in demand, as our model assumes. In the model we also assume that the losses suffered by those in declining occupations are determined in equilibrium, and if occupational labor demand is downward sloping, then an occupational labor supply response may cushion these losses. Our finding that earnings losses in declining occupations were associated with significant outflows from these occupations suggests that this mechanism may be relevant in our context.

Several of our findings are inconsistent with the predictions of the frictionless version of the model: we find that the probability of leaving declining occupations was not decreasing in initial occupational earnings rank; earnings losses due to occupational decline were decreasing (rather than increasing) in initial earnings rank; and earnings losses of those who left declining occupations were higher (rather than lower) than the losses of those who remained.

Our empirical results are more consistent with the version of the model that allows for occupational switching costs that decrease in workers' abilities in the destination occupation, since this can account for our finding that those with lower initial within-occupation earnings rank suffered larger earnings losses as a result of occupational decline.

When we allow for both differential occupational switching costs (as above) and displacement, we can account for several findings at the same time. In this case, those with lower initial within-occupation earnings rank suffer larger earnings losses as a result of occupational decline; switchers' earnings losses may be larger than those of stayers (as we find); and displacement may lead to switching probabilities that are U-shaped in initial earnings, whereby low-earning workers switch if displaced, while high-earning workers switch voluntarily.

Our empirical analysis also sheds light on the nature of the occupational switching costs in the model. In practice we find that roughly a third of the employment years lost can be accounted for by increased unemployment, and almost ten percent are due to retraining. The stronger responses to occupational decline of unemployment and retraining among lower-ranked workers further supports our interpretation of heterogeneous switching costs.

Finally, our model suggests that the effects of an adverse occupational demand shock may differ depending on whether the shock was anticipated. We find that unanticipated declines are generally associated with smaller earnings losses and smaller mobility responses. The former is consistent with our model, but the latter is not. A possible explanation may be that conditional on predicted occupational employment growth, our declining indicator isolates a lower level of exposure to actual Swedish employment declines than in the unconditional regression, or the one only conditioning on individual characteristics.⁴¹

6 Conclusion

In this paper, we study the long-run employment and earnings losses that workers suffer when demand for their occupations declines. We begin by measuring anticipated and actual occupational declines in the US, which we map into panel micro data on Swedish workers. We find that even after controlling for key predictors of occupational decline, employment changes in declining Swedish occupations were around 20-40 log points lower than in non-declining occupations.

Despite this large fall in employment, we find that over 28 years, those who in 1985 worked

⁴¹However, exposure declines by less than the mobility response, in relative terms. See columns (2) and (4) in Tables 1 and 4.

in declining occupations experienced earnings (employment) losses that were around 2-5 (1-2) percent of mean cumulative earnings (employment), compared to those who initially worked in non-declining occupations. The earnings losses are on the higher end of the above-mentioned range when we control only for individual covariates, and lower when we also control for anticipated occupational changes and industry and occupation characteristics. Around a third of the cumulative employment losses are accounted for by increased unemployment, and a further tenth by increased time spent in government retraining. Further evidence from a panel of US workers, while noisier, suggests that mean employment and earnings losses were no larger than in Sweden.

We find that workers in the bottom tercile of their occupations' earnings distributions suffered the largest losses (around 8-11 percent). Workers in the bottom tercile also lost more years of employment and spent more time in unemployment and retraining. We find that those in declining occupations were significantly more likely to leave their starting occupations. The propensity to exit declining occupations was U-shaped in initial occupational earnings rank, with those at the bottom (and to a lesser extent at the top) more likely to leave their starting occupations.

We show that our findings are consistent with a Roy model with negative occupational demand shocks, where workers may suffer displacement, and where finding reemployment takes time. In the model, those at the bottom of a declining occupation also have low earnings capacity in other occupations, and therefore find it harder to find reemployment—whether in their own occupations or in other occupations. Hence they lose most from occupational decline. The model also rationalizes the U-shaped exit pattern that we describe above: those at the bottom of their occupations' earnings distributions are more likely to leave their occupations when they are displaced, while those at the top are more likely to leave to avoid negative demand shocks.

Our findings suggest that the mean losses of occupational decline are lower than the losses suffered by displaced workers that have been reported in prior literature. This is likely because occupational decline is typically gradual, and can be partly managed through retirements, reduced entry into declining occupations, and increased job-to-job exits to other occupations.

Gradual occupational decline may also impose fewer negative spillovers on local economies compared to large, sudden shocks, such as plant closures.

At the same time, future occupational decline could still have substantial adverse consequences for workers' outcomes, for the following three reasons. First, our paper studies occupational decline that—while unanticipated early in workers' careers—was nevertheless fairly gradual. But if, for example, machine learning improves rapidly, occupational replacement may happen faster, and may be accompanied by an overall worsening of employment opportunities (Bostrom, 2014). Second, the occupational decline that we study largely spared the most skilled occupations, but this may change with new technologies. Many professionals made sizeable investments in skills that are particularly useful in their occupations, and some may also benefit from economic rents. It is possible that for these workers the earnings losses from future occupational decline may be higher than those we estimate. Finally, and perhaps most importantly, our findings show that low-earning individuals are already suffering considerable (pre-tax) earnings losses, even in Sweden, where institutions are geared towards mitigating those losses and facilitating occupational transitions. Helping these workers stay productive when they face occupational decline remains an important challenge for governments.

References


- ACEMOGLU, D., AND P. RESTREPO (2019): "Automation and New Tasks: How Technology Displaces and Reinstates Labor," Working Paper 25684, National Bureau of Economic Research.
- ADERMON, A., AND M. GUSTAVSSON (2015): "Job Polarization and Task-Biased Technological Change: Evidence from Sweden, 1975-2005," *Scandinavian Journal of Economics*, 117(3), 878–917.
- ANDERSSON, J. (2017): "Insurances against job loss and disability: Private and public interventions and their effects on job search and labor supply," Ph.D. thesis, Uppsala University.
- ARNTZ, M., T. GREGORY, AND U. ZIERAHN (2017): "Revisiting the risk of automation," *Economics Letters*, 159, 157 – 160.
- AUTOR, D., AND D. DORN (2009): "This Job Is "Getting Old": Measuring Changes in Job Opportunities Using Occupational Age Structure," *American Economic Review*, 99(2), 45–51.

- AUTOR, D. H. (2015): “Why Are There Still So Many Jobs? The History and Future of Workplace Automation,” *Journal of Economic Perspectives*, 29(3), 3–30.
- AUTOR, D. H., D. DORN, G. H. HANSON, AND J. SONG (2014): “Trade Adjustment: Worker-Level Evidence,” *The Quarterly Journal of Economics*, 129(4), 1799–1860.
- BANA, S. H. (2019): “Identifying Vulnerable Displaced Workers: The Role of State-Level Occupation Conditions,” working paper.
- BESSEN, J. (2016): “How computer automation affects occupations: Technology, jobs, and skills,” Law and economics research paper 15-49, Boston University School of Law.
- BOSTROM, N. (2014): *Superintelligence: Paths, Dangers, Strategies*. Oxford University Press.
- BRYNJOLFSSON, E., AND A. MCAFEE (2014): *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*. W. W. Norton & Company.
- BUREAU OF LABOR STATISTICS (1986): *Occupational Outlook Handbook, 1986-87 Edition: Bulletin of the United States Bureau of Labor Statistics, No. 2250*. US Department of Labor.
- (2018): *Occupational Outlook Handbook, 2018-2019 Edition*. Bernan Press.
- CAMPBELL, P. (2018): “Trucks headed for a driverless future,” *Financial Times*, January 31.
- CAPRETTINI, B., AND H.-J. VOTH (2017): “Rage Against the Machines: Labour-Saving Technology and Unrest in England, 1830-32,” CEPR Discussion Papers DP11800, Centre for Economic Policy Research.
- CASELLI, F., AND A. MANNING (2018): “Robot arithmetic: new technology and wages,” *American Economic Review: Insights*.
- CORTES, G. M. (2016): “Where Have the Middle-Wage Workers Gone? A Study of Polarization Using Panel Data,” *Journal of Labor Economics*, 34(1), 63–105.
- CORTES, G. M., AND G. GALLIPOLI (2017): “The costs of occupational mobility: An aggregate analysis,” *Journal of the European Economic Association*, 16(2), 275–315.
- DAL BO, E., F. FINAN, O. FOLKE, T. PERSSON, AND J. RICKNE (2019): “Economic Losers and Political Winners: Sweden’s Radical Right,” working paper.
- DAUTH, W., S. FINDEISEN, J. SUEDEKUM, AND N. WOESSNER (2018): “Adjusting to Robots: Worker-Level Evidence,” Working Papers 13, Federal Reserve Bank of Minneapolis, Opportunity and Inclusive Growth Institute.


- DAVIS, S., AND T. VON WACHTER (2011): “Recessions and the Costs of Job Loss,” *Brookings Papers on Economic Activity*, 42(2 (Fall)), 1–72.
- EDIN, P.-A., AND B. HOLMLUND (1995): “The Swedish Wage Structure: The Rise and Fall of Solidarity Wage Policy?,” in *Differences and Changes in Wage Structures*, NBER Chapters, pp. 307–344. National Bureau of Economic Research, Inc.
- ELIASON, M., AND D. STORRIE (2006): “Lasting or Latent Scars? Swedish Evidence on the Long-Term Effects of Job Displacement,” *Journal of Labor Economics*, 24(4), 831–856.
- FREY, C. B., AND M. A. OSBORNE (2017): “The future of employment: How susceptible are jobs to computerisation?,” *Technological Forecasting and Social Change*, 114, 254 – 280.
- GALAASEN, S. M., AND A. R. KOSTØL (2018): “Mismatch and the Consequences of Job Loss,” working paper.
- GATHMANN, C., I. HELM, AND U. SCHÖNBERG (2018): “Spillover effects in local labor markets: Evidence from mass layoffs,” *Journal of the European Economic Association*.
- GATHMANN, C., AND U. SCHÖNBERG (2010): “How general is human capital? A task-based approach,” *Journal of Labor Economics*, 28(1), 1–49.
- GOLDIN, C. (2000): “Labor Markets in the Twentieth Century,” in *The Cambridge History of the United States*, ed. by R. E. G. Stanley L. Engerman, vol. 3, pp. 549–624. Cambridge University Press.
- GOOS, M., A. MANNING, AND A. SALOMONS (2014): “Explaining Job Polarization: Routine-Biased Technological Change and Offshoring,” *American Economic Review*, 104(8), 2509–26.
- GOTTFRIES, N. (2018): “The labor market in Sweden since the 1990s,” *IZA World of Labor*, pp. 411–411.
- JACOBSON, L. S., R. J. LALONDE, AND D. G. SULLIVAN (1993): “Earnings Losses of Displaced Workers,” *The American Economic Review*, 83(4), 685–709.
- KAMBOUROV, G., AND I. MANOVSKII (2009): “Occupational specificity of human capital,” *International Economic Review*, 50(1), 63–115.
- LINDBECK, A. (1997): “The Swedish Experiment,” *Journal of Economic Literature*, 35(3), 1273–1319.
- LINDQVIST, E., AND R. VESTMAN (2011): “The Labor Market Returns to Cognitive and Noncognitive Ability: Evidence from the Swedish Enlistment,” *American Economic Journal: Applied Economics*, 3(1), 101–128.
- MARX, K. (1867): *Das Kapital*. Verlag von Otto Meisner.

- NEAL, D. (1995): “Industry-Specific Human Capital: Evidence from Displaced Workers,” *Journal of labor Economics*, 13(4), 653–677.
- OECD (2015): “Back to Work: Sweden,” Discussion paper.
- PARENT, D. (2000): “Industry-Specific Capital and the Wage Profile: Evidence from the National Longitudinal Survey of Youth and the Panel Study of Income Dynamics,” *Journal of Labor Economics*, 18(2), 306–323.
- PAVAN, R. (2011): “Career Choice and Wage Growth,” *Journal of Labor Economics*, 29(3), 549–587.
- PISSARIDES, C. A. (2000): *Equilibrium Unemployment Theory, 2nd Edition*. The MIT Press.
- POLETAEV, M., AND C. ROBINSON (2008): “Human capital specificity: evidence from the Dictionary of Occupational Titles and Displaced Worker Surveys, 1984–2000,” *Journal of Labor Economics*, 26(3), 387–420.
- ROY, A. D. (1951): “Some Thoughts on the Distribution of Earnings,” *Oxford Economic Papers*, 3(2), 135–146.
- SCHEIBER, N. (2018): “High-Skilled White-Collar Work? Machines Can Do That, Too,” *New York Times*, July 07.
- SCHMILLEN, A. (2018): “Vocational education, occupational choice and unemployment over the professional career,” *Empirical Economics*.
- SKANS, O. N., P.-A. EDIN, AND B. HOLMLUND (2009): “Wage dispersion between and within plants: Sweden 1985-2000,” in *The structure of wages: An international comparison*, pp. 217–260. University of Chicago Press.
- SPITZ-OENER, A. (2006): “Technical change, job tasks, and rising educational demands: Looking outside the wage structure,” *Journal of Labor Economics*, 24(2), 235–270.
- SULLIVAN, P. (2010): “Empirical Evidence on Occupation and Industry Specific Human Capital,” *Labour Economics*, 17(3), 567–580.
- SUSSKIND, R. E., AND D. SUSSKIND (2015): *The Future of the Professions: How Technology Will Transform the Work of Human Experts*. Oxford University Press.
- VENERI, C. M. (1997): “Evaluating the 1995 occupational employment projections,” *Monthly Labor Review*, 120.
- VIKSTRÖM, J., AND G. VAN DEN BERG (2017): “Långsiktiga effekter av arbetsmarknadsutbildning,” Discussion Paper Rapport 2017:17.


(a) No switching cost




(b) Constant switching cost c




(b') Displacement (constant cost)



(c) Heterogenous switching cost $C - \alpha_{iB}$




(c') Displacement (heter. cost)



Notes: k_{it} denotes the occupation chosen by worker i in period t . α_{ik} denotes log productivity of worker i in occupation k . d is the amount by which the relative occupational log price declines from period 1 to period 2. The parameter values chosen are $(\bar{\alpha}, d, c, \hat{c}, C) = (1, 0.5, 0.25, 0.25, 1)$.

Figure 1: Sorting in a two-period Roy model



Notes: Diamonds mark the coefficients on the declining indicator from regressions of annual earnings or cumulative earnings on the indicator, including the same controls as in columns (2) ('Individual controls') and (6) ('Occupation and industry controls') of Table 4, separately for each year 1986-2013. Bars indicate 95-percent confidence intervals.

Figure 2: Differences in earnings and cumulative earnings by exposure to occupational decline, over time

Table 1: Quantifying workers' exposure to occupational decline

	(1)	(2)	(3)	(4)	(5)	(6)
<i>A. Workers aged 16-64 in 1985 (3,061,051 observations)</i>						
Declining	-0.49 (0.12)	-0.44 (0.11)	-0.43 (0.11)	-0.31 (0.10)	-0.28 (0.11)	-0.22 (0.10)
<i>B. Workers aged 25-36 in 1985 (877,324 observations)</i>						
Declining	-0.47 (0.11)	-0.40 (0.11)	-0.39 (0.11)	-0.28 (0.10)	-0.27 (0.12)	-0.22 (0.10)
Demographics & earnings		✓	✓	✓	✓	✓
Life-cycle profiles			✓	✓	✓	✓
Predictors of growth				✓	✓	✓
Occupation dummies					✓	✓
Industry dummies						✓

Notes: Results from regressions of occupational log employment changes on a dummy for working in a declining occupation are shown. Regressions are run on individual-level data. However, the dependent variable is the difference in aggregate log employment in Swedish 3-digit occupations between 2013 and 1985, matched to each workers' 1985 5-digit occupation using a cross-walk. A Swedish 5-digit occupation is classified as 'Declining' if there are employment losses of more than 25 percent between 1986-2016 in the corresponding US occupation(s). In the regressions reported here, the 'Declining' variable indicates that an individual worked in such an occupation in 1985. Demographic controls include female, cohort, county, and education dummies. Occupation-level life-cycle profiles are cumulative earnings calculated for each individual based on their 1985 occupation. Predictors of growth include 1985 employment shares, 1960-85 occupational employment growth, and the predicted growth index. Occupation and industry dummies are at the 1-digit and 2-digit levels, respectively. Robust standard errors, clustered by 1985 3-digit occupation, in parentheses.

Table 2: Baseline characteristics of workers in subsequently declining occupations

	(1) Female	(2) Age	(3) Compulsory school	(4) High school	(5) College	(6) Earnings	(7) Manufacturing
<i>A. Workers aged 16-64</i>							
Intercept	0.52 (0.078)	39.5 (0.41)	0.33 (0.030)	0.56 (0.033)	0.11 (0.027)	191.3 (10.8)	0.25 (0.050)
Declining	-0.25 (0.088)	-0.89 (0.63)	0.13 (0.035)	-0.063 (0.034)	-0.070 (0.028)	-0.23 (11.0)	0.38 (0.085)
<i>B. Workers aged 25-36</i>							
Intercept	0.51 (0.078)	30.8 (0.078)	0.23 (0.022)	0.64 (0.033)	0.13 (0.032)	182.8 (9.28)	0.23 (0.050)
Declining	-0.26 (0.085)	-0.19 (0.091)	0.15 (0.030)	-0.065 (0.034)	-0.082 (0.034)	12.0 (9.40)	0.38 (0.084)

Notes: Results from OLS regressions of various baseline (1985) characteristics on a constant and an indicator for working in a declining occupation are shown (see the notes to Table 1 for the definition of the declining indicator). Earnings are measured in thousand Swedish crowns inflated to 2014 levels. The sample includes all individuals of the indicated ages who were employed and earned at least the base amount in 1985, and whose education, occupation, and industry are observed. The number of observations is 3,061,051 in panel A and 877,324 in panel B. Robust standard errors, clustered by 1985 3-digit occupation, in parentheses.

Table 3: Balance of pre-determined characteristics

	(1)	(2)	(3)	(4)	(5)	(6)
<i>A. Military test scores</i>						
	Cognitive skills			Non-cognitive skills		
Declining	-0.24 (0.084)	-0.015 (0.035)	0.022 (0.022)	-0.20 (0.062)	-0.077 (0.041)	-0.022 (0.021)
Individual controls		✓	✓		✓	✓
Occ. & industry controls			✓			✓
Mean of dep. var.		0.06			0.06	
Observations			272,350			
<i>B. Mother's characteristics</i>						
	Mother finished high school			Mother's earnings (1985)		
Declining	-0.059 (0.020)	-0.012 (0.0100)	0.0033 (0.0058)	-6.73 (2.41)	-2.31 (1.29)	0.079 (0.84)
Individual controls		✓	✓		✓	✓
Occ. & industry controls			✓			✓
Mean of dep. var.		0.35			97.4	
Observations			609,075			
<i>C. Father's characteristics</i>						
	Father finished high school			Father's earnings (1985)		
Declining	-0.069 (0.027)	-0.0088 (0.012)	0.0075 (0.0067)	-13.7 (6.67)	-1.38 (2.85)	2.26 (1.99)
Individual controls		✓	✓		✓	✓
Occ. & industry controls			✓			✓
Mean of dep. var.		0.43			174.2	
Observations			451,301			

Notes: Results from regressions of various pre-determined characteristics on a dummy for working in 1985 in a subsequently declining occupation are shown. Test scores from the military enlistment are standardized to have mean zero and unit variance within enlistment cohorts. The sample in panel A includes men born in Sweden from 1952-1959 with non-missing test scores (more than 85 percent of men in each cohort), who were employed and earned at least the base amount in 1985, and whose education, occupation, and industry are observed. The samples in panels B and C are the same as that in panel B of Table 2, except that individuals with missing information on mother's or father's education and income were dropped. Individual-level controls include female, cohort, county, and education dummies, as well as earnings in 1985. Occupation and industry controls include predicted life-time income, predictors of occupational growth, occupation dummies, and industry dummies. Robust standard errors, clustered by 1985 3-digit occupation, in parentheses.

Table 4: Occupational decline and individual-level cumulative employment and earnings 1986-2013

	(1)	(2)	(3)	(4)	(5)	(6)
<i>A. Cumulative years employed 1986-2013 (mean: 23.4)</i>						
Declining	-0.73 (0.26)	-0.49 (0.20)	-0.49 (0.20)	-0.30 (0.20)	-0.24 (0.18)	-0.19 (0.14)
<i>B. Cumulative real earnings ('000 2014 SEK) 1986-2013 (mean: 6,926)</i>						
Declining	-354 (419)	-347 (120)	-241 (81)	-117 (76)	-63 (71)	-126 (58)
<i>C. Cumulative real earnings divided by predicted initial earnings (mean: 38.7)</i>						
Declining	-4.29 (0.91)	-2.10 (0.53)	-2.21 (0.54)	-1.52 (0.54)	-0.98 (0.41)	-1.11 (0.36)
Demographics & earnings		✓	✓	✓	✓	✓
Life-cycle profiles			✓	✓	✓	✓
Predictors of growth				✓	✓	✓
Occupation dummies					✓	✓
Industry dummies						✓

Notes: Results from regressions of the indicated outcomes on a dummy for working in 1985 in a subsequently declining occupation are shown. Demographic controls include female, cohort, county, and education dummies. Occupation-level life-cycle profiles are cumulative earnings calculated for each individual based on their 1985 occupation. Predictors of growth include 1985 employment shares, 1960-85 occupational employment growth, and the predicted growth index. Occupation and industry dummies are at the 1-digit and 2-digit levels, respectively. The number of observations is 877,324. The sample is the same as that in panel B of Table 2. Robust standard errors, clustered by 1985 3-digit occupation, in parentheses.

Table 5: Occupational decline and individual occupational stability

	(1)	(2)	(3)	(4)	(5)	(6)
<i>A. Probability of working in same 3-digit occupation in 2013 as in 1985 (mean: 0.29)</i>						
Declining	-0.14 (0.043)	-0.11 (0.041)	-0.11 (0.042)	-0.065 (0.032)	-0.086 (0.035)	-0.045 (0.020)
<i>B. Probability of working in same 2-digit occupation in 2013 as in 1985 (mean: 0.35)</i>						
Declining	-0.12 (0.034)	-0.088 (0.034)	-0.087 (0.035)	-0.051 (0.030)	-0.070 (0.030)	-0.037 (0.019)
<i>C. Probability of working in same 1-digit occupation in 2013 as in 1985 (mean: 0.40)</i>						
Declining	-0.098 (0.030)	-0.070 (0.031)	-0.069 (0.032)	-0.039 (0.029)	-0.060 (0.027)	-0.031 (0.018)
Demographics & earnings		✓	✓	✓	✓	✓
Life-cycle profiles			✓	✓	✓	✓
Predictors of growth				✓	✓	✓
Occupation dummies					✓	✓
Industry dummies						✓

Notes: Results from regressions of the indicated outcomes on a dummy for working in 1985 in a subsequently declining occupation are shown. See the notes to Tables 1 and 4 for the definition of the declining indicator and a description of control variables, respectively. The number of observations is 553,169. The sample is the same as that in panel B of Table 2, except that individuals who were employed in 2013 but not sampled in the Wage Structure Statistics had to be excluded, as it is unknown whether they work in the same occupation in 2013 as in 1985. Sampling weights are applied. Robust standard errors, clustered by 1985 3-digit occupation, in parentheses.

Table 6: Heterogeneity by within-occupation earnings rank

	Employment (1)	(2)	Earnings (3)	(4)	Earnings, normalized (5)	(6)	Remain (7)	(8)
<i>A. Linear interaction</i>								
Declining	-0.51 (0.21)	-0.23 (0.15)	-353.5 (110.7)	-131.0 (55.8)	-2.16 (0.55)	-1.19 (0.37)	-0.11 (0.041)	-0.045 (0.020)
Declining \times rank	1.17 (0.34)	1.17 (0.30)	441.5 (142.3)	449.2 (146.8)	2.63 (0.58)	2.63 (0.57)	-0.011 (0.023)	-0.0010 (0.017)
<i>B. Dummy interactions</i>								
Declining	-0.32 (0.24)	-0.031 (0.18)	-323.2 (123.8)	-98.0 (66.7)	-1.94 (0.54)	-0.97 (0.41)	-0.083 (0.045)	-0.022 (0.021)
Declining \times bottom tercile	-1.12 (0.35)	-1.13 (0.33)	-341.8 (106.7)	-350.1 (101.5)	-2.10 (0.54)	-2.06 (0.51)	-0.046 (0.014)	-0.040 (0.013)
Declining \times top tercile	0.54 (0.20)	0.55 (0.16)	232.3 (135.8)	235.1 (132.1)	1.37 (0.43)	1.40 (0.48)	-0.047 (0.027)	-0.030 (0.018)
Individual controls	✓	✓	✓	✓	✓	✓	✓	✓
Occupation & industry controls		✓		✓		✓		✓
Mean of dep. var.		23.4		6,926		38.7		0.29
Mean of dep. var., bottom		22.3		6,001		35.6		0.27
Observations				877,324				553,786

Notes: Results from regressions of the indicated outcomes on the declining indicator, within-occupation earnings rank or tercile dummies (coefficients omitted from table), and their interactions are shown. Within-occupation earnings ranks are computed in 1985 and re-scaled so as to range from -1 to 1 . In panel A, the main effect on the declining indicator thus applies to the individual earning the median income within her occupation, and the coefficient on the interaction gives the inter-quartile range. Individual-level controls include female, cohort, county, and education dummies, as well as earnings in 1985. Occupation and industry controls include predicted life-time income, predictors of occupational growth, occupation dummies, and industry dummies. Normalized earnings are cumulative earnings divided by initial predicted earnings. The sample for columns (1)-(6) is the same as that in Table 4, and for columns (7)-(8) it is the same as that in Table 5. Robust standard errors, clustered by 1985 3-digit occupation, in parentheses.

Table 7: Occupational decline and the incidence of unemployment and retraining

	Ever			Cumulative days			
	(1)	(2)	(3)	(4)	(5)	(6)	(8)
<i>A. Unemployment</i>							
Declining	0.041 (0.021)	0.013 (0.013)	0.015 (0.012)	0.019 (0.015)	52.4 (24.8)	17.9 (14.0)	20.5 (18.2)
Declining \times rank			-0.036 (0.012)			-63.8 (21.5)	
Declining \times bottom tercile				0.017 (0.012)			42.4 (18.3)
Declining \times top tercile				-0.033 (0.012)			-43.7 (17.0)
Mean of dep. var.						262	
Mean of dep. var., bottom			0.39 0.43			317	
<i>B. Retraining</i>							
Declining	0.035 (0.010)	0.012 (0.0064)	0.013 (0.0063)	0.015 (0.0081)	11.4 (2.68)	4.73 (1.46)	5.81 (2.26)
Declining \times rank			-0.027 (0.0070)			-8.63 (1.98)	
Declining \times bottom tercile				0.014 (0.0072)			4.38 (2.28)
Declining \times top tercile				-0.022 (0.0064)			-6.96 (2.12)
Mean of dep. var.			0.13			29	
Mean of dep. var., bottom			0.15			35	
Individual controls	✓	✓	✓	✓	✓	✓	✓
Occupation & industry controls		✓	✓	✓	✓	✓	✓

Notes: Results from regressions of the indicated outcomes on the declining indicator, within-occupation earnings rank or tercile dummies (coefficients omitted from table), and their interactions are shown. Incidence of unemployment and retraining are measured during the period 1992-2013. The sample is the same as that in panel B of Table 2. See the notes to Table 6 for a description of right-hand side variables. Robust standard errors, clustered by 1985 3-digit occupation, in parentheses.

Table 8: Occupational decline and older workers

	Cumulative years employed			Cumulative earnings		Age at retirement			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
A. Workers aged 37-48 in 1985 (976,637 observations)									
Declining	-0.70 (0.16)	-0.32 (0.11)	-0.47 (0.12)	-273.1 (53.0)	-72.9 (43.4)	-99.4 (39.7)	-0.39 (0.097)	-0.15 (0.065)	-0.25 (0.074)
Declining × rank			0.98 (0.25)			173.6 (85.8)			0.65 (0.18)
Mean of dependent variable		17.2			4,759			62.8	
B. Workers aged 49-60 in 1985 (650,538 observations)									
Declining	-0.29 (0.085)	-0.047 (0.070)	-0.087 (0.072)	-75.0 (18.2)	12.3 (18.8)	8.09 (18.2)	-0.19 (0.062)	-0.011 (0.048)	-0.038 (0.049)
Declining × rank			0.18 (0.093)			14.4 (26.4)			0.13 (0.072)
Mean of dependent variable		7.0			1,576			63.6	
Individual controls	✓	✓	✓	✓	✓	✓	✓	✓	✓
Occupation & industry controls		✓	✓		✓	✓		✓	✓

Notes: Results from regressions of the indicated outcomes on the declining indicator, within-occupation earnings rank (coefficient omitted from table), and their interaction are shown. Retirement is defined as the beginning of a continuous spell of years with zero earnings lasting until age 65. Samples are as in panel A of Table 2, but restricted by age as indicated. See the notes to Table 6 for a description of right-hand side variables. Robust standard errors, clustered by 1985 3-digit occupation, in parentheses.

Online Appendix for “Individual Consequences of Occupational Decline”

Per-Anders Edin

Tiernan Evans

Georg Graetz

Sofia Hernnäs

Guy Michaels*

June 4, 2019

*Corresponding author: Guy Michaels (g.michaels@lse.ac.uk).

Contents

A	Theory appendix	3
A.1	Setting	3
A.2	Baseline model	4
A.3	Switching costs	6
A.3.1	Constant switching costs	6
A.3.2	Heterogenous switching costs	8
A.4	Displacement	9
A.4.1	Displacement under constant switching costs	9
A.4.2	Displacement under heterogenous switching costs	11
A.5	Anticipation	12
A.6	Summary	13
B	Data appendix	15
B.1	Data sources	15
B.1.1	Occupation data	15
B.1.2	Swedish microdata	15
B.1.3	NLSY data	16
B.2	Construction of variables	17
B.2.1	Occupation variables	17
B.2.2	OOH to Sweden	17
B.2.2.1	Assigning US OOH employment growth to Swedish occupations given a hypothetically unchanging OOH classification	18
B.2.2.2	Assigning US OOH employment growth to Swedish occupations given the changing OOH classification	19
B.2.2.3	Identifying technology-related declines	21
B.2.3	Swedish micro-level variables	22
B.2.4	NLSY variables	23
B.2.4.1	Occupation and industry codes	23
B.2.4.2	Income	25
B.2.4.3	Employment	27
B.2.4.4	Occupation life-cycle profiles	28
B.2.4.5	Individual controls	29
B.2.4.6	Sample weighting	29
B.3	Sample Restrictions	29
B.3.1	Swedish Data Sample restrictions	29
B.3.2	NLSY Data Sample restrictions	29
C	Appendix figures and tables	33

A Theory appendix

Here we provide a self-contained exposition of the theoretical model discussed in Section 2 of the paper, including all formal derivations. We consider two occupations, one of which is hit by a negative demand shock. We investigate how workers' likelihood of leaving the affected occupation, and their earnings losses, depend on their initial earnings. Starting from a standard frictionless Roy (1951) model, we successively introduce positive and potentially heterogeneous costs of switching occupation; as well as the possibility that workers are displaced from their jobs and incur a cost to find a new job even when remaining in their initial occupation. Finally, we consider how workers' sorting differs when the negative demand shock is anticipated.

A.1 Setting

We consider a competitive economy with a continuum of individuals indexed by i who live for two periods $t \in \{1, 2\}$ and each supplies a unit of labor inelastically each period. There are two occupations indexed by $k \in \{A, B\}$ for the workers to choose from. Workers' per-period log earnings are given by $y_{ikt} = \pi_{kt} + \alpha_{ik} - c_{ikt}$ where π_{kt} is the time-varying and stochastic (log) price of a unit of output in occupation k , α_{ik} is the time-invariant (log) amount of output that worker i produces in occupation k , and $c_{ikt} \geq 0$ is a time cost related to occupational switching, which we discuss below.¹ There are no saving opportunities and earnings are consumed immediately. We define the life-time expected utility function as $\mathbb{E}[y_{ik1} + \beta y_{ik2}]$, where $\beta > 0$ is a discount factor. In each period, workers choose the occupation that maximizes their expected utility. As a normalization, we assume that workers always choose occupation A if indifferent. Since we focus our analysis on relative wages, we define $\tilde{\pi}_t \equiv \pi_{Bt} - \pi_{At}$ and assume for simplicity that $\tilde{\pi}_1 = 0$.² Prices are determined in equilibrium by supply and demand. However, here we take them as given, and analyze the consequences of a change to prices occurring in period 2 for occupational sorting and earnings. Note that the second period may be interpreted as all periods following this change, so β could be larger than one. For simplicity, we assume that α_{iA} and α_{iB} are independent and both uniformly distributed between zero and some finite but possibly large number $\bar{\alpha}$. We explain in the following subsections that our main results are robust to alternative distributional assumptions.

In period 2, there is a negative demand shock to occupation A such that $\pi_{A2} - \pi_{A1} = -d$ and $\tilde{\pi}_2 = d, d > 0$. This may be due to labor-replacing technology becoming available, or cheaper, in occupation A . We are interested in the consequences of the shock for the earnings of workers who start out in occupation A , under various assumptions about switching costs and anticipation of the price change. Formally, let $l_i \equiv \mathbb{E}[y_{i2} | \alpha_{iA}, k_{i1} = A, D_A = 0] - \mathbb{E}[y_{i2} | \alpha_{iA}, k_{i1} = A, D_A = 1]$ be the expected earnings loss in period 2 that results from the shock, conditional on worker i starting out in occupation A , and conditional on her ability (and hence earnings rank) α_{iA} , where the occurrence of the shock is indicated by $D_A \in \{0, 1\}$. Similarly, l_i^{switch} and l_i^{stay} denote the earnings losses further conditioned on leaving and staying, respectively, and p_i

¹The time cost may reflect search or retraining (or both); we assume throughout that a worker's wage equals the value of her marginal product, $e^{\pi_{kt} + \alpha_{ik}}$. We thus abstract from any job-level rents that may arise in the presence of search frictions.

²We do not claim to identify any aggregate gains from technological change, and we do not model them here.

is the probability of switching.³ The overall loss is given by

$$l_i = l_i^{\text{stay}} - p_i (l_i^{\text{stay}} - l_i^{\text{switch}}). \quad (\text{OA1})$$

As long as there is no displacement then $l_i^{\text{stay}} = d$ and by revealed preference $l_i^{\text{switch}} \leq d$, so that $l_i \leq d$. Thus, switching enables workers to mitigate the losses from occupational decline. In the following subsections we verify that, in each version of our model, $\frac{\partial p_i}{\partial d} \geq 0$, $\frac{\partial l_i}{\partial d} \geq 0$ (with strict inequalities for some i): the larger the drop in demand, the more workers switch, and the higher are earnings losses. Furthermore, $\frac{\partial l_i}{\partial \alpha_{iA}} = -\frac{\partial p_i}{\partial \alpha_{iA}} (l_i^{\text{stay}} - l_i^{\text{switch}}) + p_i \frac{\partial l_i^{\text{switch}}}{\partial \alpha_{iA}}$. In other words, losses decrease with initial within-occupation earnings rank if the switching probability is increasing and the loss of switchers decreasing in initial earnings rank, $\frac{\partial p_i}{\partial \alpha_{iA}} > 0$ and $\frac{\partial l_i^{\text{switch}}}{\partial \alpha_{iA}} < 0$.

In what follows, we investigate how mean earnings losses vary with α_{iA} , and hence with initial earnings, under various assumptions about switching costs and anticipation of the price change. To characterize switching behavior and earnings losses, we require a distributional assumption. For simplicity, we henceforth assume that α_{iA} and α_{iB} are independent and both uniformly distributed between zero and some finite but possibly large number $\bar{\alpha}$. We argue below that our results are robust to alternative distributional assumptions.

A.2 Baseline model

We start with the simplest case, where occupational prices π_{kt} are revealed at the start of each period and there are no switching costs. Hence, occupational choice is a sequence of static decisions that can be analyzed in isolation. The set of workers choosing occupation A in period 1 is characterized by the inequality $\alpha_{iB} \leq \alpha_{iA}$, and it lies on and below the main diagonal in panel (a) of Figure OA1 (blue and red areas). The workers who switch in the second period must satisfy the inequalities $\alpha_{iB} \leq \alpha_{iA}$ and $\alpha_{iB} > \alpha_{iA} - d$, indicated by the blue area in panel (a) of Figure OA1.


To characterize switching probabilities and earnings losses, we need to distinguish two cases. Among workers in occupation A with $\alpha_{iA} \leq d$, everyone switches and their period-2 log earnings, given uniformity, are on average $\alpha_{iA}/2$, which is also the earnings loss they suffer. For those with $\alpha_{iA} > d$, the probability of switching is d/α_{iA} . The switchers' log productivity in occupation B lies between $\alpha_{iA} - d$ and α_{iA} , so given uniformity their period-2 log earnings are on average $\alpha_{iA} - d/2$, so that they suffer a loss of $d/2$. Switching probabilities, and their derivatives with respect to initial skill, are thus

$$p_i = \begin{cases} 1 & \text{if } \alpha_{iA} \leq d \\ \frac{d}{\alpha_{iA}} & \text{if } \alpha_{iA} > d, \end{cases} \quad \frac{\partial p_i}{\partial \alpha_{iA}} = \begin{cases} 0 & \text{if } \alpha_{iA} \leq d \\ -\frac{d}{\alpha_{iA}^2} < 0 & \text{if } \alpha_{iA} > d, \end{cases}$$


³Formally,

$$\begin{aligned} l_i &\equiv l_i(\alpha_{iA}, d) &\equiv \mathbb{E}[y_{i2} | \alpha_{iA}, k_{i1} = A, D_A = 0] - \mathbb{E}[y_{i2} | \alpha_{iA}, k_{i1} = A, D_A = 1], \\ l_i^{\text{switch}} &\equiv l_i^{\text{switch}}(\alpha_{iA}, d) &\equiv \mathbb{E}[y_{i2} | \alpha_{iA}, k_{i1} = A, D_A = 0] - \mathbb{E}[y_{i2} | \alpha_{iA}, k_{i1} = A, k_{i2} = B, D_A = 1], \\ l_i^{\text{stay}} &\equiv l_i^{\text{stay}}(\alpha_{iA}, d) &\equiv \mathbb{E}[y_{i2} | \alpha_{iA}, k_{i1} = A, D_A = 0] - \mathbb{E}[y_{i2} | \alpha_{iA}, k_{i1} = A, k_{i2} = A, D_A = 1], \\ p_i &\equiv p_i(\alpha_{iA}, d) &\equiv \mathbb{P}(k_{i2} = B | k_{i1} = A, \alpha_{iA}, D_A = 1). \end{aligned}$$


(a) No switching cost




(b) Constant switching cost c




(b') Displacement (constant cost)



(c) Heterogenous switching cost $C - \alpha_{iB}$



(c') Displacement (heter. cost)



Notes: k_{it} denotes the occupation chosen by worker i in period t . α_{ik} denotes log productivity of worker i in occupation k . d is the amount by which the relative occupational log price declines from period 1 to period 2. The parameter values chosen are $(\bar{\alpha}, d, c, \hat{c}, C) = (1, 0.5, 0.25, 0.25, 1)$.

Figure OA1: Sorting in a two-period Roy model

and earnings losses are

$$l_i = \begin{cases} \frac{\alpha_{iA}}{2} & \text{if } \alpha_{iA} \leq d \\ d \left(1 - \frac{d}{2\alpha_{iA}}\right) & \text{if } \alpha_{iA} > d, \end{cases} \quad \frac{\partial l_i}{\partial \alpha_{iA}} = \begin{cases} \frac{1}{2} > 0 & \text{if } \alpha_{iA} \leq d \\ \frac{d^2}{2\alpha_{iA}^2} > 0 & \text{if } \alpha_{iA} > d. \end{cases}$$

Given the above expressions, it is also straightforward to verify that $\partial p_i / \partial d \geq 0$ and $\partial l_i / \partial d \geq 0$. We summarize our analytical results for the baseline model as follows.

Result 1 *The fraction who switch among those initially working in A is weakly decreasing in α_{iA} . Moreover, switchers' earnings losses are also weakly increasing in α_{iA} ; and taken together, mean earnings losses for workers starting out in occupation A are strictly increasing in α_{iA} , and hence initial occupational earnings.*

To understand the intuition for these results, call occupation A “typist” and occupation B “cashier”, where typists suffer a negative demand shock. The worst typists could only become the worst cashiers, otherwise they would have chosen to be cashiers in period 1. But the best typists can at most become the best cashiers, and in general they will not all be the best cashiers. Therefore, the best typists are less able to mitigate their earnings losses by becoming cashiers, and they suffer larger losses than the worst typists.

This argument suggests that switching probabilities are decreasing and earnings losses are increasing in ability under a large set of alternative assumptions on the skill distributions. A sufficient condition for earnings losses to be higher for the most able than for the least able is that there is finite support with positive probability mass for all $(\alpha_{iA}, \alpha_{iB}) \in [0, \bar{\alpha}] \times [0, \bar{\alpha}]$.

A.3 Costs of switching between occupations

We continue to assume that the period-2 price change is unanticipated, but now we assume that there are costs of switching occupations. We think of these costs as the time lost searching for a new job or spent in retraining, and model them as additive in log terms. We start with the simple case where the time cost is constant across workers (and thus proportional to earnings), and then consider a case where it is decreasing in workers' ability in the destination occupation B.

A.3.1 Constant switching costs

Take first the case where the switching cost for moving from occupation A to B is a constant $c \in (0, d)$; the case $c \geq d$ is uninteresting since nobody would switch in response to the adverse shock, so we only consider the case $c < d$. Occupational choice is no longer a period-by-period decision. Instead, workers choose in period 1 the occupation with the highest expected present discounted value of log earnings, net of switching costs. Let us assume that occupational log prices follow a random walk, $\mathbb{E}[\tilde{\pi}_2] = \tilde{\pi}_1 = 0$, where the last equality is due to our earlier simplifying assumption.⁴

⁴Instead of the random walk assumption we could impose that demand changes are somehow otherwise perfectly unforeseen, for instance due to adaptive expectations (in Section A.5 we consider the case where demand changes are anticipated).

If choosing occupation A, expected life-time utility is $V_{iA} = \alpha_{iA} + \beta \mathbb{E}[\max\{\pi_{A2} + \alpha_{iA}, \pi_{B2} + \alpha_{iB} - c\}]$. If choosing occupation B, it is $V_{iB} = \alpha_{iB} + \beta \mathbb{E}[\max\{\pi_{A2} + \alpha_{iA}, \pi_{B2} + \alpha_{iB}\}]$. Consider the following exhaustive list of possible cases:

- If $\alpha_{iA} \geq \alpha_{iB}$, then $\mathbb{E}[\max\{\pi_{A2} + \alpha_{iA}, \pi_{B2} + \alpha_{iB} - c\}] = \mathbb{E}[\max\{\pi_{A2} + \alpha_{iA}, \pi_{B2} + \alpha_{iB}\}] = \mathbb{E}[\pi_{A2}] + \alpha_{iA}$, and $V_{iA} \geq V_{iB}$.
- If $\alpha_{iB} - c \leq \alpha_{iA} < \alpha_{iB}$, then $\mathbb{E}[\max\{\pi_{A2} + \alpha_{iA}, \pi_{B2} + \alpha_{iB} - c\}] = \mathbb{E}[\pi_{A2}] + \alpha_{iA}$, $\mathbb{E}[\max\{\pi_{A2} + \alpha_{iA}, \pi_{B2} + \alpha_{iB}\}] = \mathbb{E}[\pi_{B2}] + \alpha_{iB}$, and $V_{iA} < V_{iB}$.
- If $\alpha_{iA} < \alpha_{iB} - c$, then $\mathbb{E}[\max\{\pi_{A2} + \alpha_{iA}, \pi_{B2} + \alpha_{iB} - c\}] = \mathbb{E}[\pi_{B2}] + \alpha_{iB} - c$, $\mathbb{E}[\max\{\pi_{A2} + \alpha_{iA}, \pi_{B2} + \alpha_{iB}\}] = \mathbb{E}[\pi_{B2}] + \alpha_{iB}$, and $V_{iA} < V_{iB}$.

This establishes that worker i chooses occupation A in period 1 if and only if $\alpha_{iA} \geq \alpha_{iB}$. Decisions at the beginning of the terminal period 2 are easily characterized, as before. After the price change, worker i switches if and only if $\alpha_{iB} - c > \alpha_{iA} - d$. In sum, the workers who switch to occupation B after the price change satisfy the inequalities $\alpha_{iB} \leq \alpha_{iA}$ and $\alpha_{iB} > \alpha_{iA} - (d - c)$. Panel (b) of Figure OA1 shows a situation that is qualitatively similar to the baseline model, except that the blue region marking the workers who switch is smaller than in panel (a). Deriving expressions for the switching probability and earnings loss as a function of initial earnings follows along very similar lines as in the proof of Result 1.

To characterize switching probabilities and earnings losses, we again distinguish two cases. Among workers in occupation A with $\alpha_{iA} \leq d - c$, everyone switches and their period-2 log earnings, given uniformity, are on average $\alpha_{iA}/2 - c$, which gives an average loss of $\alpha_{iA}/2 + c$. For those with $\alpha_{iA} > d - c$, the probability of switching is $(d - c)/\alpha_{iA}$. The switchers' log productivity in occupation B lies between $\alpha_{iA} - (d - c)$ and α_{iA} , so given uniformity their period-2 log earnings are on average $\alpha_{iA} - (d - c)/2$, so that they suffer a loss of $(d + c)/2$. Switching probabilities, and their derivatives with respect to initial skill, are thus

$$p_i = \begin{cases} 1 & \text{if } \alpha_{iA} \leq d - c \\ \frac{d-c}{\alpha_{iA}} & \text{if } \alpha_{iA} > d - c, \end{cases} \quad \frac{\partial p_i}{\partial \alpha_{iA}} = \begin{cases} 0 & \text{if } \alpha_{iA} \leq d - c \\ -\frac{d-c}{\alpha_{iA}^2} < 0 & \text{if } \alpha_{iA} > d - c, \end{cases}$$

and earnings losses are

$$l_i = \begin{cases} \frac{\alpha_{iA}}{2} + c & \text{if } \alpha_{iA} \leq d - c \\ d - \frac{(d-c)^2}{2\alpha_{iA}} & \text{if } \alpha_{iA} > d - c, \end{cases} \quad \frac{\partial l_i}{\partial \alpha_{iA}} = \begin{cases} \frac{1}{2} > 0 & \text{if } \alpha_{iA} \leq d - c \\ \frac{(d-c)^2}{2\alpha_{iA}^2} > 0 & \text{if } \alpha_{iA} > d - c. \end{cases}$$

Given the above expressions, as in the baseline model, $\partial p_i / \partial d \geq 0$ and $\partial l_i / \partial d \geq 0$. We summarize our analytical results for the constant switching cost model as follows.

Result 2 *Under a constant switching cost, we obtain the same qualitative conclusions as in Result 1: The fraction who switch among those initially working in A is weakly decreasing in α_{iA} . Moreover, switchers' earnings losses are also weakly increasing in α_{iA} ; and taken together, mean earnings losses for workers starting out in occupation A are strictly increasing in α_{iA} , and hence initial occupational earnings.*

The same intuition as in the baseline model of Section A.2 applies: the best workers in the declining occupation are less likely to be able to mitigate their earnings losses by switching occupation.

A.3.2 Heterogenous switching costs

Suppose instead that workers who wish to switch from A to B must pay a switching cost equal to $C - \alpha_{iB}$, with $C > \bar{\alpha}$ (the condition $C > \bar{\alpha}$ ensures that all workers face a strictly positive switching cost). This structure of switching costs captures in a reduced form way the frictions that occupational moves may entail: for example, job search may take time, and those more able in the new occupation may find a job more quickly.

We continue to assume that occupational log prices follow a random walk. If choosing occupation A , we have $V_{iA} = \alpha_{iA} + \beta \mathbb{E}[\max\{\pi_{A2} + \alpha_{iA}, \pi_{B2} + \alpha_{iB} - (C - \alpha_{iB})\}]$. If choosing occupation B , then $V_{iB} = \alpha_{iB} + \beta \mathbb{E}[\max\{\pi_{A2} + \alpha_{iA}, \pi_{B2} + \alpha_{iB}\}]$. Consider the following exhaustive list of possible cases.

- If $\alpha_{iA} \geq \alpha_{iB}$, then $\mathbb{E}[\max\{\pi_{A2} + \alpha_{iA}, \pi_{B2} + \alpha_{iB} - (C - \alpha_{iB})\}] = \mathbb{E}[\max\{\pi_{A2} + \alpha_{iA}, \pi_{B2} + \alpha_{iB}\}] = \mathbb{E}[\pi_{A2}] + \alpha_{iA}$, and $V_{iA} \geq V_{iB}$.
- If $\alpha_{iB} - (C - \alpha_{iB}) \leq \alpha_{iA} < \alpha_{iB}$, then $\mathbb{E}[\max\{\pi_{A2} + \alpha_{iA}, \pi_{B2} + \alpha_{iB} - (C - \alpha_{iB})\}] = \mathbb{E}[\pi_{A2}] + \alpha_{iA}$, $\mathbb{E}[\max\{\pi_{A2} + \alpha_{iA}, \pi_{B2} + \alpha_{iB}\}] = \mathbb{E}[\pi_{B2}] + \alpha_{iB}$, and $V_{iA} < V_{iB}$.
- If $\alpha_{iA} < \alpha_{iB} - (C - \alpha_{iB})$, then $\mathbb{E}[\max\{\pi_{A2} + \alpha_{iA}, \pi_{B2} + \alpha_{iB} - (C - \alpha_{iB})\}] = \mathbb{E}[\pi_{B2}] + \alpha_{iB} - (C - \alpha_{iB})$, $\mathbb{E}[\max\{\pi_{A2} + \alpha_{iA}, \pi_{B2} + \alpha_{iB}\}] = \mathbb{E}[\pi_{B2}] + \alpha_{iB}$, and $V_{iA} < V_{iB}$.

This establishes again that worker i chooses occupation A in period 1 if and only if $\alpha_{iA} \geq \alpha_{iB}$. After the price change in period 2, worker i switches if and only if $\alpha_{iB} - (C - \alpha_{iB}) > \alpha_{iA} - d$. Thus, the workers who switch to occupation B after the shock must now satisfy the inequalities $\alpha_{iB} \leq \alpha_{iA}$ and $\alpha_{iB} > \alpha_{iA}/2 + (C - d)/2$, shown as the blue area in panel (c) of Figure OA1. The figure shows that workers with α_{iA} below $C - d$ do not switch, and that above $C - d$, the fraction switching is increasing in α_{iA} due to uniformity. Thus,

$$p_i = \begin{cases} 0 & \text{if } \alpha_{iA} < C - d \\ \frac{1}{2} - \frac{C-d}{2\alpha_{iA}} & \text{if } \alpha_{iA} \geq C - d, \end{cases} \quad \frac{\partial p_i}{\partial \alpha_{iA}} = \begin{cases} 0 & \text{if } \alpha_{iA} < C - d \\ \frac{C-d}{2\alpha_{iA}^2} > 0 & \text{if } \alpha_{iA} \geq C - d, \end{cases}$$

and

$$l_i^{\text{switch}} = \frac{d}{2} + \frac{C - \alpha_{iA}}{2}, \quad \frac{\partial l_i^{\text{switch}}}{\partial \alpha_{iA}} = -\frac{1}{2} < 0$$

where we used the fact that mean earnings of switchers equal $3\alpha_{iA}/2 - (C + d)/2$. Thus, if $\alpha_{iA} > C - d$, we have by (OA1) that $\partial l_i / \partial \alpha_{iA} < 0$ (and zero otherwise). It is also straightforward to verify that $\partial p_i / \partial d \geq 0, \partial l_i / \partial d > 0$. To summarize:

Result 3 *If the cost of switching occupations is decreasing in α_{iA} , then the fraction who switch among those initially working in A is weakly increasing in α_{iA} , and mean losses conditional on α_{iA} are (weakly) decreasing in α_{iA} , and hence initial earnings.*

In terms of the example above, in this case the worst typists do not switch, because their initial choice of occupation A reveals not only low earnings potential in occupation B but also a large switching cost. Among the best typists, however, many possess substantial earnings potential as cashiers, as well as low switching costs. Therefore, the best typists are on average better able to mitigate their earnings losses by

becoming cashiers, and hence the earnings losses from the demand shock are smaller for the best typists than for the worst typists.⁵

A.4 Job displacement

So far, we have been concerned with earnings losses as a function of initial earnings in the context of a Roy model where any moves between occupations are voluntary. By revealed preference, losses of movers must be less than those of stayers. Here we show that introducing job displacement and a cost of finding a new job in the initial occupation may overturn this result.⁶

Suppose that workers who start in occupation A experience job displacement with probability λ at the end of period 1. For simplicity, and to maximize similarity with previous cases, we assume that displacement catches workers by surprise: ex-ante, they believe the probability of displacement equals zero. We have verified that our results are qualitatively unchanged when we assume that workers know the true probability before choosing an occupation in period 1.

A.4.1 Displacement under constant switching costs

Displacement affects choices only in the presence of switching costs. First we assume that displaced workers incur a cost $\hat{c} > 0$ to find a job in A , and a cost c to find a job in B (the latter of course also applies to non-displaced workers). Here we have in mind exogenous job losses, for instance due to plant closure, which are a standard feature of search models (see for instance Pissarides, 2000).

The workers who are displaced switch occupation if and only if $\alpha_{iB} > \alpha_{iA} - (d - (c - \hat{c}))$, and among them are individuals who would remain if not displaced, $\alpha_{iB} \leq \alpha_{iA} - (d - c)$. Workers not suffering displacement switch voluntarily if and only if $\alpha_{iB} > \alpha_{iA} - (d - c)$. Thus, there is a set of workers who switch occupation only if suffering displacement, as illustrated by the yellow area in panel (b') of Figure 1. Given uniformity, we have switching probabilities

$$p_i = \begin{cases} 1 & \text{if } \alpha_{iA} \leq d - c \\ \lambda + (1 - \lambda) \frac{d - c}{\alpha_{iA}} & \text{if } d - c < \alpha_{iA} < d - (c - \hat{c}) \\ \frac{d - c}{\alpha_{iA}} + \lambda \frac{\hat{c}}{\alpha_{iA}} & \text{if } \alpha_{iA} \geq d - (c - \hat{c}), \end{cases} \quad \frac{\partial p_i}{\partial \alpha_{iA}} = \begin{cases} 0 & \text{if } \alpha_{iA} \leq d - c \\ -(1 - \lambda) \frac{d - c}{\alpha_{iA}^2} < 0 & \text{if } d - c < \alpha_{iA} < d - (c - \hat{c}) \\ -\frac{d - c}{\alpha_{iA}^2} - \lambda \frac{\hat{c}}{\alpha_{iA}^2} < 0 & \text{if } \alpha_{iA} \geq d - (c - \hat{c}). \end{cases}$$

The earnings loss in the first region is the same as in the case without displacement in Section A.3.1. The calculations are more involved in the second and third regions. Let ρ_i^{blue} denote the probability that a worker's skill α_{iB} lies in the blue region in panel (b') of Figure OA1 (conditional on starting out in occupation A and on α_{iA}), let l_i^{blue} be her expected loss, and analogously define ρ_i^{yellow} , ρ_i^{red} and

⁵While our model excludes occupation-specific human capital, it does allow us to think about some of its potential implications. For example, if all workers accumulate occupation-specific human capital additively (in logarithms) the effects are similar to adding constant switching costs, since switching means foregoing this capital. And if higher ability workers accumulate more occupation-specific human capital they become less occupationally mobile, in contrast to the case of heterogeneous switching costs discussed above.

⁶Recall that a large literature has documented substantial earnings losses due to job displacement (see for instance Jacobson, LaLonde, and Sullivan, 1993) and even larger losses if such displacement coincides with switching occupation (Kambourov and Manovskii, 2009).

$l_i^{\text{yellow}}, l_i^{\text{red}}$. When $d - c < \alpha_{iA} < d - (c - \hat{c})$, we have

$$\left(\rho_i^{\text{yellow}}, \rho_i^{\text{blue}}\right) = \left(1 - \frac{d - c}{\alpha_{iA}}, \frac{d - c}{\alpha_{iA}}\right), \quad \left(l_i^{\text{yellow}}, l_i^{\text{blue}}\right) = \left(\lambda \left(\frac{\alpha_{iA}}{2} + \frac{d + c}{2}\right) + (1 - \lambda)d, \frac{d + c}{2}\right)$$

and

$$l_i = l_i^{\text{blue}} + \rho_i^{\text{yellow}} \left(l_i^{\text{yellow}} - l_i^{\text{blue}}\right),$$

where we note that $l_i^{\text{yellow}} - l_i^{\text{blue}} = \lambda \frac{\alpha_{iA}}{2} + (1 - \lambda) \frac{d - c}{2} > 0$.

When $\alpha_{iA} \geq d - (c - \hat{c})$, we additionally define ρ_i^{red} and l_i^{red} in the same sense as above, so we have

$$\left(\rho_i^{\text{red}}, \rho_i^{\text{yellow}}, \rho_i^{\text{blue}}\right) = \left(1 - \frac{d - (c - \hat{c})}{\alpha_{iA}}, \frac{\hat{c}}{\alpha_{iA}}, \frac{d - c}{\alpha_{iA}}\right), \quad \left(l_i^{\text{red}}, l_i^{\text{yellow}}, l_i^{\text{blue}}\right) = \left(\lambda(d + \hat{c}) + (1 - \lambda)d, d + \frac{\hat{c}}{2}, \frac{d + c}{2}\right)$$

and

$$l_i = l_i^{\text{blue}} + \rho_i^{\text{red}} \left(l_i^{\text{red}} - l_i^{\text{blue}}\right) + \rho_i^{\text{yellow}} \left(l_i^{\text{yellow}} - l_i^{\text{blue}}\right),$$

where $l_i^{\text{red}} - l_i^{\text{blue}} = \frac{d - c}{2} + \lambda \hat{c} > 0$ and $l_i^{\text{yellow}} - l_i^{\text{blue}} = \frac{d - c}{2} + \frac{\hat{c}}{2} > 0$. Thus,

$$l_i = \begin{cases} \frac{\alpha_{iA}}{2} + c & \text{if } \alpha_{iA} \leq d - c \\ \frac{d + c}{2} + \left(1 - \frac{d - c}{\alpha_{iA}}\right) \left(\lambda \frac{\alpha_{iA}}{2} + (1 - \lambda) \frac{d - c}{2}\right) & \text{if } d - c < \alpha_{iA} < d - (c - \hat{c}) \\ d + \lambda \hat{c} - \frac{(d - (c - \hat{c}))(d - c + 2\lambda \hat{c})}{2\alpha_{iA}} & \text{if } \alpha_{iA} \geq d - (c - \hat{c}), \end{cases}$$

and

$$\frac{\partial l_i}{\partial \alpha_{iA}} = \begin{cases} \frac{1}{2} > 0 & \text{if } \alpha_{iA} \leq d - c \\ \frac{d - c}{\alpha_{iA}^2} \left(l_i^{\text{yellow}} - l_i^{\text{blue}}\right) + \rho_i^{\text{yellow}} \frac{\lambda}{2} > 0 & \text{if } d - c < \alpha_{iA} < d - (c - \hat{c}) \\ \frac{(d - (c - \hat{c}))(d - c + 2\lambda \hat{c})}{2\alpha_{iA}^2} > 0 & \text{if } \alpha_{iA} \geq d - (c - \hat{c}). \end{cases}$$

It is also straightforward to verify that $\partial p_i / \partial d \geq 0$ and $\partial l_i / \partial d \geq 0$. To summarize:

Result 4 *If the cost of switching occupations is constant, if workers in occupation A may be displaced from their jobs, and if the cost of finding a new job in occupation A is also constant, then the fraction who switch among those initially working in (weakly) decreasing in α_{iA} , and mean losses conditional on α_{iA} are strictly increasing in α_{iA} and hence initial earnings.*

Displacement does not affect the qualitative results that we obtained before when assuming constant switching costs (Section A.3.1). It is still the case that the most skilled among the workers starting out in A are less likely to be in a position where they can mitigate their earnings losses by switching occupation.

A.4.2 Displacement under heterogenous switching costs

In the case of a heterogeneous cost $(C - \alpha_{iB})$ of moving to occupation B , we introduce in symmetric fashion a cost of finding a job in occupation A in case of displacement, $C - \alpha_{iA}$. Period-1 occupational choices are the same as in the case of heterogenous switching costs without displacement, since we assume that workers believe the displacement probability to be zero. Recall that workers not affected by displacement switch voluntarily if and only if $\alpha_{iB} > \alpha_{iA}/2 + (C - d)/2$. Workers that do suffer displacement switch occupation if and only if $\alpha_{iB} > \alpha_{iA} - d/2$, as illustrated by the yellow area in (c') of Figure 1. The figures suggests that, going from low to high values of α_{iA} , the switching probability first equals λ as all displaced workers switch and then falls below λ and becomes a decreasing function of α_{iA} , and possibly, eventually an increasing function of α_{iA} , since the incidence of voluntary switching increases in α_{iA} for large values of α_{iA} .

Indeed, switching probabilities are characterized as

$$p_i = \begin{cases} \lambda & \text{if } \alpha_{iA} \leq \frac{d}{2} \\ \frac{\lambda d}{2\alpha_{iA}} & \text{if } \frac{d}{2} < \alpha_{iA} < C - d \\ \frac{1-\lambda}{2} - \frac{(1-\lambda)C-d}{2\alpha_{iA}} & \text{if } \alpha_{iA} \geq C - d, \end{cases} \quad \frac{\partial p_i}{\partial \alpha_{iA}} = \begin{cases} 0 & \text{if } \alpha_{iA} \leq \frac{d}{2} \\ -\frac{\lambda d}{2\alpha_{iA}^2} < 0 & \text{if } \frac{d}{2} < \alpha_{iA} < C - d \\ \frac{(1-\lambda)C-d}{2\alpha_{iA}^2} \geq 0 \Leftrightarrow (1-\lambda)C \geq d & \text{if } \alpha_{iA} \geq C - d. \end{cases}$$

The expression for the switching probability for $\alpha_{iA} \geq C - d$ follows from the fact that when $\alpha_{iA} \geq C - d$, the probability of a worker being in the blue region in panel (c') of Figure OA1 equals $1/2 - (C - d)/(2\alpha_{iA})$, and the probability of being in the yellow region is $C/(2\alpha_{iA}) - 1/2$. We see that for $\alpha_{iA} > C - d$ the switching probability is strictly increasing in α_{iA} provided the shock is not too large, $d/C < 1 - \lambda$.

The losses from occupational decline in this version of the model are

$$l_i = \begin{cases} \lambda C + (1 - \lambda)d & \text{if } \alpha_{iA} \leq \frac{d}{2} \\ \lambda \left[C - \alpha_{iA} + d - \frac{d^2}{4\alpha_{iA}} \right] + (1 - \lambda)d & \text{if } \frac{d}{2} < \alpha_{iA} < C - d \\ \ell_i & \text{if } \alpha_{iA} \geq C - d, \end{cases}$$

where ℓ_i is to be characterized below. For $\alpha_{iA} < d/2$, workers move if and only if they are displaced, and the expected loss of movers is constant at C , while the expected loss of stayers is constant at d . Hence, the overall expected loss in this region equals $\lambda C + (1 - \lambda)d$ and thus does not depend on α_{iA} . For $\alpha_{iA} \in (d/2, C - d)$, the expected loss for those who are displaced and stay is $C - \alpha_{iA} + d$ and for those who are displaced and move it is $C - \alpha_{iA} + d/2$ (all non-displaced workers stay). Together with the switching probability, the result follows.

Finally, we consider the case $\alpha_{iA} \geq C - d$. Let ρ_i^{red} denote the probability that a worker's skill α_{iB} lies in the red region in panel (c') of Figure OA1 (conditional on starting out in occupation A and on α_{iA}), let l_i^{red} be her expected loss, and analogously define ρ_i^{yellow} , ρ_i^{blue} and l_i^{yellow} , l_i^{blue} . We have

$$\ell_i = l_i^{\text{blue}} + \rho_i^{\text{red}} (l_i^{\text{red}} - l_i^{\text{blue}}) + \rho_i^{\text{yellow}} (l_i^{\text{yellow}} - l_i^{\text{blue}}),$$

where

$$(\rho_i^{\text{red}}, \rho_i^{\text{yellow}}) = \left(1 - \frac{d}{2\alpha_{iA}}, \frac{C - \alpha_{iA}}{2\alpha_{iA}}\right), \quad (l_i^{\text{red}}, l_i^{\text{yellow}}, l_i^{\text{blue}}) = \left(d + \lambda(C - \alpha_{iA}), d + \lambda \frac{C - \alpha_{iA}}{2}, \frac{d}{2} + \frac{C - \alpha_{iA}}{2}\right),$$

and so

$$\ell_i = d + \lambda(C - \alpha_{iA}) + \frac{1}{2\alpha_{iA}} \left(-\frac{d^2}{2} + (1 - \lambda)(C - \alpha_{iA}) \left(d - \frac{C - \alpha_{iA}}{2} \right) \right)$$

and

$$\frac{\partial \ell_i}{\partial \alpha_{iA}} = \frac{d^2 + C(C - 2d)(1 - \lambda) - \alpha_{iA}^2(3\lambda + 1)}{4\alpha_{iA}}.$$

Thus,⁷

$$\frac{\partial \ell_i}{\partial \alpha_{iA}} = \begin{cases} 0 & \text{if } \alpha_{iA} \leq \frac{d}{2} \\ \lambda \left[\frac{d^2}{4\alpha_{iA}^2} - 1 \right] < 0 & \text{if } \frac{d}{2} < \alpha_{iA} < C - d \\ \partial \ell_i / \partial \alpha_{iA} < 0 \iff \frac{d}{C} < \frac{2}{3} & \text{if } \alpha_{iA} \geq C - d. \end{cases}$$

Let us assume that $d/C < \min\{2/3, 1 - \lambda\}$. This is a sensible assumption, as it implies that the shock is not huge—it does not come close to making the declining occupation vanish (recall that $C > \bar{\alpha}$). This also is the case that gives rise to the richest patterns of switching behavior, as for instance the yellow region of panel (c') of Figure OA1 would not coexist with the blue region if $d/2 > C - d$. As before, $\partial p_i / \partial d \geq 0$ and $\partial l_i / \partial d \geq 0$. To summarize:

Result 5 *If the cost of switching occupations is decreasing in initial earnings, if workers in occupation A may be displaced from their jobs, and if the cost of finding a new job in occupation A is also decreasing in initial earnings, then the fraction who switch among those initially working in A is U-shaped in α_{iA} , and mean losses conditional on α_{iA} are (weakly) decreasing in α_{iA} and hence initial earnings.*

Intuitively, the earnings loss is decreasing in initial earnings, as in the case with heterogeneous switching costs without displacement (Section A.3.2), because the costs of moving jobs—both within and across occupations—decrease with initial earnings.

A.5 Revelation of period-2 prices at the start of period 1

As a final variation on our model, we consider a case where period-2 prices are revealed to be $\tilde{\pi}_2 = d$ at the start of period 1. Without switching costs, decisions are again static and occupational choices follow the same conditions as in the baseline model of Section A.2. Suppose however that there is a constant (across individuals) switching cost $c \in (0, d)$ for moving from A to B, as in the first scenario considered

⁷The sign of $\frac{\partial \ell_i}{\partial \alpha_{iA}}$ is the same as the sign of its numerator. Since $\alpha_{iA} > C - d$ when $l_i = \ell_i$, the numerator is strictly less than

$$d^2 + C(C - 2d)(1 - \lambda) - (C - d)^2(3\lambda + 1) = \lambda \left(d^2 - 4(C - d)^2 \right)$$

which is negative if $d/C < 2/3$. The expression is also negative if $d/C > 2$, but cases with $d > C$ are uninteresting since they imply that everyone leaves the declining occupation.

in Section A.3. Workers choose an occupational path by comparing the deterministic life-time utilities associated with the choices (A,A) , (A,B) , and (B,B) . Let switching costs again be constant. The life-time utilities are given by $V_{iAA} = \alpha_{iA} + \beta(\alpha_{iA} - d)$, $V_{iAB} = \alpha_{iA} + \beta(\alpha_{iB} - c)$, and $V_{iBB} = \alpha_{iB} + \beta\alpha_{iB}$. First, let us assume that switching costs are not too large, $(1 + \beta)c < d$. Then we have:

- If $\alpha_{iB} \leq \alpha_{iA} - (d - c)$, the worker chooses (A,A) .
- If $\alpha_{iB} > \alpha_{iA} - (d - c)$ and $\alpha_{iB} \leq \alpha_{iA} - \beta c$, the worker chooses (A,B) .
- If $\alpha_{iB} > \alpha_{iA} - \beta c$, the worker chooses (B,B) .

All workers with $\alpha_{iB} > \alpha_{iA} - \beta c$ choose occupation B in period 1 and remain there. Thus, some workers who otherwise would have started out in occupation A instead start in B to avoid the switching cost, and the fraction of workers switching in the period when the shock hits is smaller than without anticipation of the shock.

If switching costs are large instead, $(1 + \beta)c \geq d$, then workers with $\alpha_{iB} \leq \alpha_{iA} - \beta c$ choose (A,A) and workers with $\alpha_{iB} > \alpha_{iA} - \beta c$ choose (B,B) , so that no switching occurs after period 1. To summarize:

Result 6 *If period-2 prices are revealed already at the start of period 1, and under a constant occupational switching cost, the fraction of workers starting out in occupation A, and the fraction of workers leaving occupation A after the first period, are both smaller than in the case without anticipation discussed in Section A.3.1. The fraction switching occupation after period 1 may even be zero if the switching cost is large.*

More generally, the model suggests that the set of workers who are in declining occupations may differ for anticipated and unanticipated shocks. Different combinations of anticipation, general equilibrium responses, heterogeneity of occupational switching costs, and displacement, may lead to a range of different outcomes.

A.6 Summary of theoretical results

We have modelled occupational decline using a Roy model, where employment in an occupation declines as a result of a fall in occupational price caused by a technology shock. The model illustrates how earnings losses due to occupational decline are mitigated by occupational switching.

Furthermore, our frictionless baseline model makes three predictions: the probability of leaving a declining occupation is decreasing in initial earnings; earnings losses due to occupational decline are increasing in initial earnings; and earnings losses of those who leave a declining occupation are less than the losses of those who remain.

Anticipating that these predictions are inconsistent with our empirical findings, we have considered several modifications to the model. Introducing an occupational switching cost that is decreasing in the worker's earnings in the destination occupation leads to a positive relationship between switching probabilities and initial earnings, and a negative relationship between earnings losses and initial earnings. Allowing for displacement, together with a cost of switching jobs within an occupation, implies that switchers' earnings losses may be larger than those of stayers. Moreover, displacement can cause

switching probabilities to be U-shaped in initial earnings, whereby low-earning workers switch involuntarily if displaced, while high-earning workers switch voluntarily regardless of displacement.

The importance of switching costs in our theoretical analysis motivates our empirical approach of focusing not only on losses in career earnings incurred by workers starting out in declining occupations, but also on losses in years employed, as well as on the incidence of unemployment and retraining. While our model does not include a non-work sector, it could be shown that a negative demand shock would trigger moves from the affected occupation into non-participation.

Finally, we have used our model to show that much of re-sorting in response to a technology shock may occur before the shock hits if it is anticipated in advance, motivating our investigation of both anticipated and unanticipated occupational decline.

B Data appendix

B.1 Data sources

Our main analysis is based on individual-level longitudinal administrative data covering the entire population of Sweden 1985-2013, and on various editions of the Occupational Outlook Handbook (OOH) published by the Bureau of Labor Statistics (BLS) in the US. Part of our analysis also uses data from the National Longitudinal Survey of Youth (NLSY) containing a sample of US residents, as well as a number of auxiliary data sources, as described below.

B.1.1 Occupation data

Our primary source for measuring occupational decline are the 1986-87 and the 2018-19 editions of the Occupational Outlook Handbook (Bureau of Labor Statistics, 1986, 2018d). The OOH contains a description of the nature of work, the current number of jobs, and projected employment growth for hundreds of occupations. For a subset of these occupations, more detailed information is reported, including required qualifications, pay, and the role of technology: whether technology is expected to affect—or has already affected—the occupation in question, and if so, what the impact on employment will be or has been. In the 1986-87 edition, 401 occupations are described, covering about 80 percent of US employment. Detailed information is available for 196 occupations, covering about 60 percent of employment.⁸

B.1.2 Swedish microdata

The main outcomes we study come from Swedish microdata. We obtain basic demographic (year of birth, gender, education, and county of residence) and labor market (employment status, annual earnings, and industry) variables from the Integrated Database for Labour Market Research (LISA), a collection of administrative registers that is—like all our other Swedish data sources—provided by Statistics Sweden. During the period 1985-2013, LISA contains one observation per year on every individual aged 16-64 living in Sweden. Employment status and industry (as well as county of residence) are measured in November each year.

We also use individual-level data from the Swedish Public Employment Service (PES), which contain information on the total number of days registered with the PES, number of days registered as unemployed, and number of days spent in retraining programs administered by the PES, for all individuals ever registered with the PES during the years 1992-2013.

Our data on workers' occupations come from the population censuses, which were conducted every five years from 1960-1990, and from the Wage Structure Statistics (WSS) for the years 1996-2013.⁹ The WSS contains the population of public sector workers and a sample of about 50 percent of private sector workers. Sampling is at the level of firms, and large firms are over-sampled. We apply sampling weights when working with the occupation variable from the WSS.

⁸The number of distinct occupations in the OOH, as well as the number of occupations covered in detail, tends to increase over time. This means that our crosswalk from the 1986-87 to the 2018-19 edition is mostly, though not always, one-to-many.

⁹We also obtain individual-level earnings data for 1975 and 1980 from the population censuses, which we use for falsification checks.

A useful feature of our data is that, in the 1985 and 1990 censuses, workers' occupation is coded using a 5-digit classification, YRKE5, containing about 1,400 distinct occupations. This allows us to accurately merge occupation-level information from the US (see below). Unfortunately, such detailed occupation codes are not available after 1990. From 1996-2013, a 3-digit classification containing 172 distinct codes, SSYK96, is available in the WSS. This classification is of a different nature than YRKE5, and the cross-walk between YRKE5 and SSYK96 likely introduces measurement error in workers' occupations after 1990.¹⁰ This is an important caveat to our analysis of occupational employment shifts and individual workers' occupational mobility during 1985-2013.

Finally, adding the 1960 census allows us to calculate prior occupational employment changes at the 3-digit level using the YRKE3 classification, a coarser version of YRKE5 (there are 229 distinct codes that cover the period 1960-85).¹¹

B.1.3 NLSY data

Since the main focus of our study is Sweden, which has better data, we try wherever possible to select and analyze US data in a way that is as close as possible to what we do in Sweden. We only depart from this when data availability or quality necessitate using alternative approaches.

The main dataset we use to study occupational decline and its consequences in the US is the National Longitudinal Survey of Youth (NLSY 1979), because it is one of the few panel datasets that are representative of a relevant age group in the US during the period we want to study. NLSY79 has a detailed set of occupation codes that are important for our analysis, since they can be readily matched to the 1986 Occupational Outlook Handbook (OOH).

Specifically, for years through 2010 we use the 1979–2010 release, from Böhm (2013)¹² with updated weights to include only those in the sample as of 1987 (see below), and updates for recent errata from the NLSY.¹³ For 2012 and 2014, we use the 1979–2014 data release.

The NLSY79 Cohort is comprised of individuals born between 1957 and 1964. These people were beginning their careers in the late 1980s, the time of interest identified in the analysis of Swedish data. NLSY79 surveys were conducted annually from 1979–1994 and on a biennial basis thereafter. We use data until and including the 2014 round, which covers earnings until 2013—the year in which our Swedish data end.

To construct the occupational life-cycle profiles, discussed in detail below, we need a larger sample than is available in the NLSY. As in Acemoglu and Autor (2011), we use individual-level data containing information on age, gender, race, education, employment status, occupation, hours and weeks worked, as well as annual labor income from the 1980 US Census, accessed through the IPUMS website (Ruggles, Alexander, Genadek, Goeken, Schroeder, and Sobek, 2010). We construct education and income variables in the same way as Acemoglu and Autor (2011).

To convert income to \$2014, we use the chained Consumer Price Index for All Urban Consumers

¹⁰Within broad types of jobs, SSYK96 also distinguishes occupations by the skill level of the workers.

¹¹The Swedish word *yrke* means occupation. SSYK stands for (the Swedish translation of) Swedish Standard Classification of Occupations.

¹²We thank Michael Böhm for generously sharing this data and his expertise.

¹³We use updated income for the “Revised Income Variables Incorrectly Coded” (Bureau of Labor Statistics, 2018b) and updated occupations for “Erroneous Occupation Codes (2002 and 2004)” (Bureau of Labor Statistics, 2018a)

(CPI-U), published by the BLS and made available by the Federal Reserve Bank of Minneapolis (2018).

B.2 Construction of variables

B.2.1 *Occupation variables (OOH)*

Using the reported employment numbers from our two editions of the OOH, we calculate the percentage growth in employment 1984-2016.¹⁴ We manually map occupations across the two editions. If, after a careful search, a 1986-87 occupation has no counterpart in the 2018-19 edition, we classify it as having vanished, and assigned a percentage growth of -100.¹⁵ While few occupations actually disappeared, examples of occupations that declined sharply include both white-collar occupations (typists, drafters, and telephone operators) and blue-collar ones (precision assemblers, welders, and butchers).

We also record for each US occupation its projected employment growth from the 1986-87 OOH. The BLS bases these predictions on (forecasts of) the size and demographic composition of the labor force, aggregate economic growth, commodity final demand, industry-level output and employment, the input-output matrix, and occupational employment and vacancies. The forecasts are not reported in percentage terms but grouped into the categories “declining”, “little or no change”, “increasing slower than average”, “increasing about as fast as average”, and “increasing faster than average”. We create a cardinal predicted growth index assigning these categories the numbers 1-5 (where higher numbers correspond to more positive predicted employment changes). We report results both from using this index and using the categorical outlook variable.¹⁶

B.2.2 *Merging of OOH variables to Swedish microdata, and defining occupational decline*

In order to merge the OOH-based variables to Swedish data, we map the 401 1986-87 OOH occupations to the 1,396 5-digit Swedish occupation codes available in the 1985 census. We successfully map 379 US occupations to 1,094 Swedish occupations—we are able to find corresponding US occupations for 91 percent of Swedish workers in 1985. We map percentage changes in US employment 1984-2016, as well as 1986-87 OOH predictions (categorical and index), to Swedish 5-digit occupations using our crosswalk, applying weights (OOH 1984 employment shares) in the case of many-to-one matches.

We define a Swedish 5-digit occupation as declining if the weighted employment growth of its corresponding OOH occupations is negative and larger (in absolute magnitude) than 25 percent. We regard this as a sensible threshold: smaller declines may be the result of measurement error, as we had to exercise judgment in matching OOH occupations over time. At the same time, we report robustness checks using a number of alternative thresholds. We also use information from the OOH to determine whether technology likely played a role in the decline.¹⁷ In 1985, 13 percent of Swedish employees worked in

¹⁴The 1986-87 OOH reports employment numbers for 1984, while the 2018-19 edition reports 2016 employment figures.

¹⁵Between the 1986-87 and 2018-19 editions of the OOH, some occupations were split or merged, which we take into account when computing the percentage growth. The details of this calculation are discussed later in this section.

¹⁶Veneri (1997) uses US employment data to evaluate the ex-post accuracy of the projections used in the 1986-87 OOH, and concludes that they correctly foresaw most occupational trends, although there were also non-trivial sources of error.

¹⁷To determine whether technology played a role in the decline, we proceeded as follows. We first applied the 25-percent cutoff to the OOH data to identify the declining occupations in the US. For the declining occupations we searched their detailed descriptions in the 1986-87 OOH for discussions of potential replacement of human labor by specific technologies, such as computers or robots. For the occupations lacking detailed descriptions in the 1986-87 OOH, we further searched one and two decades ahead, using the 1996-97 and 2006-07 editions (Bureau of Labor Statistics, 1996, 2006), since in some cases

subsequently declining occupations, and 8 percent worked in subsequently declining occupations where the decline is likely linked to technology. We now provide more details on the process.

B.2.2.1 Assigning US OOH employment growth to Swedish occupations given a hypothetically unchanging OOH classification

For clarity, we first describe what the calculation of employment growth would be if the OOH classification had not changed between the 1986-87 and 2018-19 editions. We then describe the adjustments we make given that the OOH classification did change.¹⁸

The percentage change that we assign to each Swedish occupation s in the hypothetical case of an unchanging OOH classification is given by

$$g_s \equiv \frac{N_{s,2016} - N_{s,1984}}{N_{s,1984}}, \quad (\text{OA2})$$

where $N_{s,t} \equiv \sum_{k \in \mathbb{K}_s} N_{k,t}$ is the sum of all year- t employment in the $k \in \mathbb{K}_s$ OOH occupations to which the Swedish YRKE5 occupation is matched. This percentage change can alternatively be expressed as

$$g_s \equiv \underset{1 \times K}{\alpha_s} \times \underset{K \times 1}{g}, \quad (\text{OA3})$$

where the vector α_s is a vector of weights of length K , where K is the total number of OOH occupations in the 1986-87 OOH. Each element $\alpha_{s,k}$ represents the share of OOH occupation k in the mapping to Swedish occupation s , and it is based on the employment figures in the initial period 1984.¹⁹ Thus, $\alpha_{s,k} \in [0, 1]$, the vector α_s differs between Swedish YRKE5 occupations and its elements always sum to one. The vector g is filled with the 1984-2016 growth rates of all K OOH occupations. Formally,

$$\alpha_{s,k} \equiv \frac{\mathbb{1}_{k \in \mathbb{K}_s} \times N_{k,1984}}{\sum_{k \in \mathbb{K}_s} N_{k,1984}}, \quad g_k \equiv \frac{N_{k,2016} - N_{k,1984}}{N_{k,1984}}.$$

occupations were re-grouped and so received detailed descriptions in those editions. Note that, while the OOH contains little backward-looking information on technology's role, it provides rich information on imminent technological changes expected to affect occupations. Conditional on an OOH occupation being classified as declining, we regard this information as reliable with respect to technology's role in the decline.

For those OOH occupations that we identified to have undergone technology-related declines, we map employment growth to Swedish 5-digit occupations creating a separate variable, technology-related employment growth. We define a Swedish 5-digit occupation as declining and linked to technology if the technology-related employment growth in the corresponding OOH occupations is below negative 25 percent. All technology-related declining occupations are declining occupations by construction, but some declining occupations may not be classified as having a technology link.

¹⁸In the analysis of the NLSY data, we assign the percentage change to the relevant NLSY occupational codes using the same procedure.

¹⁹Note that the 1986-87 OOH uses data from 1984. Thus, the initial period is 1984 as far as US employment figures are concerned, but the data are extracted from a 1986 publication.

The equivalence of (OA2) and (OA3) is easily shown:

$$\begin{aligned}
g_s &\equiv \frac{N_{s,2016} - N_{s,1984}}{N_{s,1984}} \\
&\equiv \frac{\sum_{k \in \mathbb{K}_s} N_{k,2016} - \sum_{k \in \mathbb{K}_s} N_{k,1984}}{\sum_{k \in \mathbb{K}_s} N_{k,1984}} \\
&\equiv \frac{\sum_{k \in \mathbb{K}_s} N_{k,1984} \times g_k}{\sum_{k \in \mathbb{K}_s} N_{k,1984}} \\
&\equiv \sum_k \frac{\mathbb{1}_{k \in \mathbb{K}_s} \times N_{k,1984} \times g_k}{\sum_{k \in \mathbb{K}_s} N_{k,1984}} \\
&\equiv \sum_k \alpha_{s,k} \times g_k \\
&\equiv \alpha_s \times g.
\end{aligned}$$

B.2.2.2 Assigning US OOH employment growth to Swedish occupations given the changing OOH classification

The computation of the total changes in equation (OA2), or the weights and changes in equation (OA3) would be straightforward if the OOH occupation classification remained constant between the 1986-87 and 2018-19 editions. Alas, it did not, and so we need to adjust the calculation for any splits and merges that took place.

To see this, consider the following example: the OOH occupation “343 Metal pourers and casters, basic shapes” had employment 12,000 in 1984. By 2016, it had been merged with sixteen other occupations to “Metal and Plastic Machine Workers”, with employment 1,039,600. It is obviously wrong to calculate the change in occupation “343 Metal pourers and casters, basic shapes” as a more than 85-fold increase:

$$g_{343} = \frac{1,039,600 - 12,000}{12,000} = 85.63$$

Instead, it is reasonable to sum the employment of all the seventeen merged occupations in 1984, with a total employment of 1,457,000, and calculate the change as

$$\hat{g}_{343} = \frac{1,039,600 - 1,457,000}{1,457,000} = -0.286$$

obtaining a 28.6% decline.

However, what happens to the weights in α_s ? If we were to weight the “343 Metal pourers and casters, basic shapes” by their adjusted employment figure for 1984 (1,457,000), this occupation would seem 121 larger than it actually was (12,000). This creates problems when “343 Metal pourers and casters, basic shapes” is matched to a Swedish YRKE5 occupation that is also matched to other OOH occupations.

Consider, for instance, the Swedish YRKE5 occupation “732.50 Precision founder” to which “343 Metal pourers and casters, basic shapes” is matched, together with another OOH occupation “344 Mold-

ers and casters, hand”.

Swedish YRKE5 occupation	OOH occupation	Employment in 1984	\hat{g}_k
732.50 Precision founder	343 Metal pourers and casters, basic shapes	12,000	-0.286
	344 Molders and casters, hand	17,000	-1.000

“344 Molders and casters, hand” was larger than “343 Metal pourers and casters, basic shapes” in 1984, and disappeared completely between 1984 and 2016. It seems like we should assign the Swedish YRKE5 occupation “732.50 Precision founders” a decline somewhere in between -28.6% and -100%, but closer to -100% since the disappearing occupation dominates. However, if we were to use *adjusted* employment figures when calculating the weights, “343 Metal pourers and casters, basic shapes” would be weighted as follows:

$$\hat{\alpha}_{s,343} = \frac{1,457,000}{1,457,000 + 17,000} = 0.988$$

That is, “343 Metal pourers and casters, basic shapes” would seem to account for almost *all* employment in the Swedish YRKE5 occupation, instead of less than half. This means that the weighted change will be mistakenly computed as

$$\begin{aligned} & \hat{\alpha}_{s,343} \times \hat{g}_{343} + \alpha_{s,344} \times \hat{g}_{344} \\ &= 0.988 \times (-0.286) + 0.012 \times (-1.00) = -0.295 \end{aligned}$$

Instead, we ought to use the original employment figures when calculating the weights. Then,

$$\alpha_{s,343} = \frac{12,000}{12,000 + 17,000} = 0.414$$

i.e. the OOH occupation “343 Metal pourers and casters, basic shapes” makes up 41.4% of employment in the Swedish YRKE5 occupation. Thus,

$$\begin{aligned} & \alpha_{s,343} \times \hat{g}_{343} + \alpha_{s,344} \times \hat{g}_{344} \\ &= 0.414 \times (-0.286) + 0.586 \times (-1.00) = -0.704 \end{aligned}$$

That is, the employment growth assigned to “732.50 Precision founders” should be -70.5%. We will thus treat weights and growth rates separately: The weights α_s are computed using the original employment figures, and the growth rates g_k are computed using the adjusted employment figures,

$$\hat{g}_s = \alpha_s \times \hat{g}. \quad (\text{OA4})$$

The formal definition of our declining indicator is thus

$$\text{Declining}_s \equiv \mathbb{1}\{\alpha_s \times \hat{g} < -0.25\}.$$

It remains to specify how exactly the growth rates should be adjusted for splits and merges.²⁰

- One-to-one: OOH occupations that were neither split or merged between the 1986-87 and 2018-19 editions of the OOH. No adjustment is needed, and the growth rate is defined as above,

$$\hat{g}_k = g_k \equiv \frac{N_{k,2016} - N_{k,1984}}{N_{k,1984}}.$$

- Many-to-one merge: Many 1984 occupations $k \in \mathbb{K}$ (where \mathbb{K} is a set of 1984 occupations) were merged into one 2016 occupation \tilde{k} . 1984 employment figures of all merged occupations are summed and compared to the 2016 figures.

$$\hat{g}_{k \in \mathbb{K}} = \frac{N_{\tilde{k},2016} - \sum_{k \in \mathbb{K}} N_{k,1984}}{\sum_{k \in \mathbb{K}} N_{k,1984}}$$

- One-to-many split: One 1984 occupation k was split into many 2016 occupations $\tilde{k} \in \tilde{\mathbb{K}}$ (where $\tilde{\mathbb{K}}$ is a set of 2016 occupations). The 2016 employment figures of all resulting splits are added and compared to the 1984 figures.

$$\hat{g}_k = \frac{\sum_{\tilde{k} \in \tilde{\mathbb{K}}} N_{\tilde{k},2016} - N_{k,1984}}{N_{k,1984}}$$

- Many-to-many: Many 1984 occupations $k \in \mathbb{K}$ (where \mathbb{K} is a set of 1984 occupations) were distributed into many 2016 occupations $\tilde{k} \in \tilde{\mathbb{K}}$ (where $\tilde{\mathbb{K}}$ is a set of 2016 occupations). The 1984 and 2016 employment figures are added and compared.

$$\hat{g}_{k \in \mathbb{K}} = \frac{\sum_{\tilde{k} \in \tilde{\mathbb{K}}} N_{\tilde{k},2016} - \sum_{k \in \mathbb{K}} N_{k,1984}}{\sum_{k \in \mathbb{K}} N_{k,1984}}$$

B.2.2.3 Identifying technology-related declines

Having calculated the adjusted employment growth \hat{g}_k for all occupations present in the 1986-87 OOH, we concentrate on those that declined sharply, $\hat{g}_k < -0.25$, and check whether there is a probable technological driver behind the decline. For this we first consult the 1986-87 OOH, and if we find nothing there, we check in the 1996 OOH (BLS, 1996), and if we still find nothing, we check the 2006 version (BLS, 2006).²¹ Each OOH occupation thus is assigned an indicator variable for technological-related decline, which equals zero whenever $\hat{g}_k \geq -0.25$, and may equal zero or one when $\hat{g}_k < -0.25$.

²⁰We have excluded four OOH occupations that were merged with or split into an unknown number of occupations: “71 Electroencephalographic technologists and technicians”, “203 Public administration — chief executives, legislators, and general administrators”, “226 Customer service representatives, utilities” and “293 Electric meter installers and repairers”.

²¹There were four heavily declining ($\hat{g}_k < -0.25$) OOH occupations where we found no information in the OOH editions of 1986, 1996, or 2006, but we still suspected technologically-related decline. Therefore, we searched in other editions of the OOH and other sources, and found potential technological drivers of occupational decline:

We can then decompose the employment growth assigned to each Swedish YRKE5 occupation as follows:

$$\hat{g}_s \equiv \underbrace{\alpha_s}_{1 \times K} \times \underbrace{\mathbb{1}\{\text{technology}\}}_{K \times K} \times \underbrace{\hat{g}}_{K \times 1} + \underbrace{\alpha_s}_{1 \times K} \times \underbrace{(I - \mathbb{1}\{\text{technology}\})}_{K \times K} \times \underbrace{\hat{g}}_{K \times 1}, \quad (\text{OA5})$$

where $\mathbb{1}\{\text{technology}\}$ is a diagonal matrix with the indicator for technologically-related decline on the diagonal, and I is the identity matrix. We define a Swedish YRKE5 as having undergone technology-related decline if it is classified as declining and if the first component of the decomposition (OA5) is less than -0.25 , formally

$$[\text{Declining (technology)}]_s \equiv \mathbb{1}\{\alpha_s \times \hat{g} < -0.25 \text{ and } \alpha_s \times \mathbb{1}\{\text{technology}\} \times \hat{g} < -0.25\}.$$

B.2.3 Swedish micro-level variables

In addition to the occupational data, we construct several variables that characterize workers' career outcomes spanning the years 1986-2013; that is, starting with the first year after we measure treatment and ending with the last year available in our data. We start by simply summing up years observed as employed and real annual labor earnings, obtaining the variables cumulative years employed and cumulative earnings.²² Following Autor, Dorn, Hanson, and Song (2014), we also create a normalized measure of cumulative earnings, whereby we divide cumulative earnings by predicted initial earnings. Cumulative earnings normalized in this way thus give the multiple of (predicted) initial earnings that a worker receives during 1986-2013.²³ We consider further earnings measures—such as rank, discounted cumulative earnings, and earnings growth—in robustness checks.

Our measure of long-run occupational mobility is a dummy variable equaling one if the individual worked in the same 3-digit SSYK96 occupation in 2013 as 1985. It equals zero if the individual works in a different occupation or is not employed.²⁴ Using the PES data, we calculate cumulative days spent unemployed and cumulative days spent in retraining during 1992-2013. We define dummy variables

213 Radio operators	“Laborsaving [sic] technical advances such as computer-controlled programming and remotely controlled transmitters” (regarding Broadcast and sound engineering technicians and radio operators, BLS 2004:260)
254 Telegraph and teletype operators	Automatic routing of calls, voice message systems (regarding Telephone operators, BLS 1994:291)
346 Motion picture projectionists	Digital projection (Hess, 2014)
391 Service station attendants	Self-service pumps at petrol stations (Emek Basker and Klimek, 2015)

²²We define a worker as employed in a given year if they are identified as working in November (when employment status is measured for the purposes of LISA) of that year and if annual earnings during that year are no lower than the base amount. When we do not observe an individual in a given year—due to emigration or death—we set employment and earnings to zero.

²³The prediction comes from a regression of log earnings on a quartic in age and dummies for gender, county, and seven education categories, run separately for each 3-digit SSYK96 occupation in 1985. We divide by predicted rather than actual initial earnings to eliminate transitory earnings variation, which would introduce an important role for mean reversion into the distribution of normalized cumulative earnings. Autor, Dorn, Hanson, and Song (2014) divide cumulative earnings by earnings averaged across four pre-treatment years for the same reason. Since we do not have annual earnings information prior to 1985, we normalize by predicted earnings instead.

²⁴Our measure of occupational mobility does not capture any temporary exits during the intervening years if workers returned to their initial occupation. A limitation of our data is that they are not conducive to studying high-frequency occupational mobility: During the years 1986-1989 and 1991-1995, we do not observe workers' occupation. And during 1996-2004, the SSYK96 variable contains substantially fewer distinct codes than from 2005 onwards.

for ever unemployed and ever having participated in retraining. As the PES data are not available for 1986-1991, we cannot capture any unemployment or retraining in these early years of our sample period. Finally, we calculate the retirement age, where we define retirement as a continuous spell of zero annual earnings up to and including age 64.²⁵

B.2.4 NLSY variables

Here we provide more detail about the construction of the variables in the NLSY data; where possible we tried to follow the procedure we used for the Swedish data, but some data constraints required us to adapt the procedure as follows.

B.2.4.1 Occupation and industry codes

As in our analysis of changes in Sweden, we use the OOH as the source for occupational employment growth and to identify declining occupations, again defining decline as a contraction in OOH-equivalent occupational employment by more than 25 percent from 1984-2016. To calculate OOH-equivalent employment growth for each occupation in the NLSY, we employ the exact same procedure as for the Swedish data and as described in the previous appendix section, with one exception as described below.

B.2.4.1.1 NLSY79 occupation data

We consider only the primary employer in our analysis and use the 1980 census code data (which is only available for the primary employer) from 1987–2000. The primary employer is determined based in CPS criteria²⁶ from 1982-1994 and is coded for each person’s main job (“job #1”) from 1994-2000 (Bureau of Labor Statistics, 2018c). From 2002 onward, NLSY79 occupation is reported only on the basis of the 2000 census codes, for all employers; we consider only the occupation associated with each person’s main job.

B.2.4.1.2 1980 to 2000 Census mapping

Because 1980 census and 2000 census occupations are not reported simultaneously, it is necessary to bridge the two. To do this, we use the tables from Autor and Dorn (2013), which convert each of the 1980 and 2000 census codes to a unique 1990 occupation code (henceforth, these unique codes are referred to as “1990 occupation code(s)”).

²⁵The LISA database includes individuals older than 64 only during later years. As we do not consistently observe individuals beyond age 64, we assume for all years that individuals aged 65 or older have retired.

²⁶The CPS employer is identified as follows:

- For those not at work during the survey week but who worked for pay since the last interview—the CPS employer is the most recent employer
- For those who worked during the survey week: for one employer—the CPS employer is the current employer; for two or more employers—the CPS employer is the employer for whom the respondent worked the most hours; for two or more employers with the same number of hours each employer—the CPS employer is the employer for whom the respondent worked the longest
- For those absent from their regular job during the survey week but who were working temporarily for another employer—the CPS employer is the current employer not the employer of absence

B.2.4.1.3 OOH 1986-87 to 2018-19

With one exception, the occupational decline calculations are identical to those used in the analysis of the Swedish data. The exception is that an additional mapping is necessary because in the NLSY data, managers are often not separated by the types of occupations they manage.

Both the 1986-87 and 2018-19 OOH include aggregate measures for some occupational groupings. In particular, for the “Managers and Administrators” grouping. We take the following steps to determine occupational growth and predicted growth for these occupations. We first separate those occupations with an exact three-way-match between the 1986-87 OOH, 2018-19 OOH, and 1980 Census codes²⁷ and calculate occupational growth for each of these occupations. We then subtract the occupational employment for these managers from the total for all managers and administrators in each of 1986-87 and 2018-19. We use these totals to calculate occupational growth for a constructed occupation: “All other managers”, which is used in the same way as any other occupation for all managers not in the three categories with an exact match.

B.2.4.1.4 Census (1980) codes to 1986-87 OOH

The OOH reports occupations on a different basis than the 1980 census occupational coding, which is used in both the NLSY79 in 1987 the 1980 census. To determine which individuals were working in declining occupations as of 1987, we create a crosswalk from the 1980 census occupations to those reported in the 1986 OOH.

We map the 1980 census codes to the 1986-87 OOH occupations primarily based on occupation description. Additionally, both the 1980 census classification of occupations and 1986 OOH classification were developed to be consistent with the 1980 Standard Occupational Classification (SOC) Manual. The major occupation groups between the two are therefore similar and also informed the mappings.

As the OOH does not cover all occupations (Bureau of Labor Statistics, 1986), there are some occupations reported in the NLSY79 that cannot be matched to OOH occupations. Because reliable data on the growth of the occupation is not available, individuals in those occupations as of 1987 are excluded from our analysis.

B.2.4.1.5 Occupation and industry groupings

For the purposes of occupation switching (see below), we group occupations based on their 1990 occupation codes, using modified groupings from Autor and Dorn (2013). Here, Autor and Dorn (2013) separately classify detailed occupations for low-skill service occupations and for non-service occupations. The only cases in which occupations are classified in both categories is as police/fire occupations also being classified as protective service. We use the protective service categorization, which also includes guards, as the relevant group for these occupations.

In addition to the occupation groups based on the 1990 occupation codes, we also group occupations based on the 1980 census codes to create base year statistics. Here, occupations are grouped based on

²⁷These occupations are (in 2018-19 OOH parlance): Education administrators, Medical and health services managers, and Property, real estate, and community association managers.

the separations (bolded and/or italicized breaks in the text) in the “1980 Occupational Codes” section of United States Bureau of the Census (1980).

Finally, we include an industry group dummy in the regressions. For this purpose, industries were grouped based on separations of the “1980 Industry Codes” section of United States Bureau of the Census (1980).

B.2.4.1.6 Occupation switching

A respondent is considered to have remained in the same occupation if the 1990 occupation code for their occupation in the year of interest is the same as the 1990 occupation code for their occupation in the initial year. We restrict the sample for comparison in each year to those interviewed in that year. Respondents whose occupation was not reported are treated as switching occupations.

Because both the occupation group and major occupation group categories are calculated from the same 1990 census codes, the comparison is for occupational category switching is direct.

B.2.4.2 Income

The income measure we use is total income from wages and salary in the prior calendar year. Reported income is truncated for privacy reasons. The procedure used in NLSY79 for top coding takes the top two percent of respondents with valid values and averages them. That averaged value replaces the values for all respondents in the top range.

We use two main measures of income in our analysis of NLSY79 outcomes. In addition to cumulative income, which is an outcome of interest in the Swedish analysis, we also consider average income.

In all cases, we follow the logic outlined in Dahl and Lochner (2012) and restrict those included in the regression to those with at least 8 years of income data. This minimum ensures that income is available throughout the period of interest rather than select years in the beginning or end. In practice, enforcing the minimum means excluding 877 person-years of reported income from 283 individuals from the average income regressions. 60% of these years of excluded reported income occurred during or before 1991; individuals with income only in the first few years cannot be reasonably compared with individuals with income throughout the period.

We also tested the sensitivity of the results to changes in the minimum number of years needed for income projection and inclusion in the regression and found no substantial difference with the results reported.

The first income measure, cumulative income, allows for direct comparability with the Swedish results. Cumulative income is calculated beginning with the 1988 income (from the 1989 survey round). Income accumulates through the last year of projection and includes years in which income was imputed or projected as outlined below.

Because cumulative income is a key variable of interest, and NLSY79 interviews are conducted only in even years beginning in 1994, we impute income where possible so that results on cumulative income, the same basis as the Swedish results, can be calculated. This procedure, which is described in detail below, is also used to impute income in survey years for individuals who were not interviewed or were missing income data, to maintain a sample of people that is as representative as possible of the US population of the relevant age groups.

Due to attrition in the sample, accumulating income over the full period reduces the size of the available sample, which is not fully resolved by the extrapolation and imputation. To more fully use the data available, we also consider average income, and compare results under alternative calculation specifications.

In the calculation most similar to that for cumulative income, average income is calculated over the years in which income is interpolated and projected, beginning with 1988 income.

To test the sensitivity of the results to our income imputation procedure, we also use an additional calculation of average income using only reported income. However, due to the missing survey years later in the sample period, this places twice the weight on early-career years (prior to 2002) as it does on mid-late career years.

To address the uneven weighting between early and late-career earnings, we add our final specification. Here, we use only the income reported in even survey years (beginning in 1990), which results in even weighting across the full period.

B.2.4.2.1 Imputation procedure

The imputation procedure largely follows that laid out in the appendix to Dahl and Lochner (2012), relying on additional information available from the NLSY to improve the imputation in the case of respondents who were deceased. We also use information for the non-survey years later in the sample to a greater extent than Dahl and Lochner (2012). We therefore treat these missing years the same as any other years in which an income was not reported.

There are 6,679 individuals in the sample who both have all necessary covariates and meet the criteria below for income imputation. Considering only these individuals, imputing and projecting income allows us to increase the weighted person-years included in the average income regression by 73%, 96% of which comes from interpolation, not extrapolation. The vast majority of this increase comes from the non-survey years, with imputation and extrapolation for these years accounting for 91% of all imputations and extrapolations. Unweighted people-years increase from 95,631 to 167,132, with nearly identical sources of the increase.

Again considering this set of eligible respondents, the number of people not responding to the NLSY increases over time. In 2000, 690 of the 6,679 included in our restricted sample were not interviewed. By 2014, that number had increased to 1,422.

Considering only the years in which a survey was conducted, although the number missing income data is growing with time, because we limit extrapolation (see below), the number of people for whom income is imputed in a survey year peaks (at 902) in 2002. In the survey years, the total number of imputed and extrapolated income data is 8,002: 64% (5,089) of these occur due to non-response to the survey; 26% (2,051) due to an unknown income; 10% (779) due to refusal; the remainder (83) due to the question being skipped.

To impute income, we used the following steps:

1. Convert income from \$1979 to \$2014, using CPI-U (Federal Reserve Bank of Minneapolis, 2018)
2. Following Dahl and Lochner (2012), restrict the analysis to individuals with income data in at least 8 years from 1986 to 2013.

3. Regress income on age and age squared (as of the middle of the year in which income is earned), using OLS separately for each individual. The income considered in the regression is all income reported for 1986 to 2013. Years in this range where no survey was conducted or the respondent wasn't surveyed or didn't answer the income question are treated as missing data in the regression.
4. For years where income data are available we use them; when they are not we use predicted income values from the regression above, adjusted as explained below.
5. To be consistent with the way NLS report income, and to avoid implausible negative labor income, we winsorize predicted values:
 - (a) To 0 at bottom end.
 - (b) To average of top 2% if in top 2%.²⁸
6. Use winsorized predicted values when income is missing and Reason for Non-Interview (RNI) is not death.
7. Set income to 0 if RNI is death, or was death in any previous year in the case of odd years where no survey was conducted.

B.2.4.2.2 Extrapolation

Extrapolation (both forward from the last year observed and backward from the first year observed), in the cases where it is used, is limited to 2 years. For example, if someone's last observed income is 2009 income (observed in 2010 survey round), predicted income is used for 2010 and 2011. Income thereafter is treated as missing. We chose to limit extrapolation to two years to strike a balance between two competing objectives: getting as many person-years as possible to keep the sample as representative as possible and not relying on the functional form of the regressions to impute values many years away from where we observe actual earnings. Extrapolating more than two years may result in implausibly large earnings if, say, the second order polynomial has a positive and increasing slope at the ends of the sample years for which we have data.

Income of 0 due to the respondent being deceased is considered known, not extrapolated, income (i.e. it is not excluded due to the cap on projections). However, in the specifications below that refer to "reported income", only the values in survey years are considered in the averages. This is done so the potential years of income are the same whether or not the respondent died during the period considered.

B.2.4.3 Employment

Data on employment and unemployment come from the weekly status arrays, which are based on the respondent's full employment history. The employment history is constructed using job tenures at each of the employers reported to the NLSY, and therefore also includes information for years in which the respondent was not interviewed. This allows us to view a more complete employment history than what is reported in the survey-year job data.

²⁸To use imputed income consistently with reported income, the same procedure was followed. If imputed or extrapolated income in a year was at least as high as the bottom threshold of the top 2% in that year, we assigned to the individual the mean of the top 2% earnings in that year.

As in the income regressions, we restrict the sample to those with at least 8 years (measured as 418 weeks) where labor force status is reported, beginning in 1987. The average weeks reported (beginning in 1988) for those included in the regressions is 1,310. As the result of the restriction on the sample, we exclude 247 individuals, with an average of 234 weeks of reported labor force status (beginning in 1988) from these regressions.

A respondent is considered employed if their reported status is employed, “associated with an employer” or in active military service. A respondent is considered unemployed if their reported status is unemployed. The remaining categories “not working (unemployment vs. out of the labor force cannot be determined)” and “out of the labor force” are considered to be out of the labor force, completing the mapping.

B.2.4.4 Occupation life-cycle profiles

To construct occupation life-cycle profiles, we require a large sample to determine how income in each occupation develops over the course of a person’s career. The NLSY does not provide a sufficient sample size for this, so we instead construct these profiles using data from the 1980 census, which uses the same occupation codes as the 1987 survey year of the NLSY79. The calculation methodology follows that used for the Swedish data. The lifecycle information of the individual occupation is used unless there were fewer than 500 people in the occupation in the 1980 census, in which case the profile for the occupation group (based on 1980 census codes) is used. The process used is outlined below.

1. In the census data, restricting the population to those of working age who have non-zero earnings, hours, and weeks worked:
 - (a) Convert pre-tax wage and salary income (“income”) from \$1979 to \$2014
 - (b) Separately for each occupation, regress log income on a quartic in age and dummies for sex, county, and education. Here we use a quartic regression because we are considering the progression of income over the full lifecycle of the occupation and have a large enough sample size to alleviate concerns of overfitting.
2. In the NLSY data
 - (a) Regress log of base year (1987) income on a quadratic in age and dummies for sex, region, education, and occupation to get predicted base year income. As the age range of this sample is much narrower, the benefits of a higher-order polynomial in age are reduced, while the smaller sample of the NLSY for a particular occupation heightens concerns of overfitting.
 - (b) Calculate predicted log base year income for each person using the fitted values from the regression above
 - (c) Generate predicted log income growth in each year by summing mean real wage growth and the expected growth based on aging, calculated by applying the occupation-specific coefficients on the quartic in age from the census regression to the change in each of those values in each year
 - (d) Calculate predicted income in each year by adding predicted log income growth to the predicted log base year income and exponentiating

Cumulative predicted income is the sum of predicted income from 1988 – 2013.

B.2.4.5 Individual controls

The full set of individual controls includes birth year, sex, region (referred to collectively as “demographics”), education, and base year income. We use the four regions available in the NSLY79 data to control for geographic variation, as state-level data is not available in the public-use NLSY79 data. For education, we use 5 categories, ranging from < High school to \geq Masters, rather than the 7 used in the Swedish analysis, as compulsory education requirements vary by state.

B.2.4.6 Sample weighting

We use the NLSY79 Custom Weighting Program to calculate weights for all individuals in the selected sample who were interviewed in the 1987 survey. Cross-sectional weights available directly from the NLS at the time of writing are incorrect due to the exclusion of 401 NLSY79 respondents from the sample when calculating the weights (Bureau of Labor Statistics, 2018a).

B.3 Sample Restrictions

B.3.1 Swedish Data Sample restrictions

Our starting sample contains all individuals born between 1921-1969—hence aged 16-64 (at some point) in 1985—who were employed in November 1985, whose annual earnings in 1985 were no less than the “base amount” (Swedish: *basbelopp*) specified by the social security administration, and about whom we have complete demographic (including education) and labor market information (including industry and occupation). The base amount is used as an accounting unit when calculating benefits, and it is typically equal to about three months’ worth of full-time work at the median wage. As we do not observe hours worked or fulltime status, we use the base amount to exclude individuals with little labor market attachment. There are 3,061,051 individuals fulfilling the above criteria.²⁹ Our *baseline sample* further restricts birth year to 1949-1960, or ages 25-36 in 1985. We drop younger workers as these are less likely to be attached to the labor market and may not yet have settled on an occupation. And we drop middle-aged and older workers from our baseline sample because we want to focus on workers who did not reach retirement age by 2013, the end of our period of study, in our main analysis. We will analyze workers born before 1949 separately.

B.3.2 NLSY Data Sample restrictions

We want to study people who have likely completed their schooling before the start of the period; at the same time, we want to use the same variation from the OOH that we used in Sweden. The balance of these two factors leads us to choose 1987 as a base year for the NLSY analysis, since by that year the youngest people covered in the NLSY will have reached age 22, and (in most cases) will have completed their education.

²⁹There were 5,281,382 individuals aged 16-64 in Sweden in 1985. Of those, 4,186,512 were employed in November 1985, and among them, 3,648,034 earned no less than the base amount during 1985. The reduction to 3,061,051 is due to missing education, industry, or occupation information, including cases where YRKE5 occupations do not have matches in the OOH.

Given the choices above, we focus on samples of people whose histories we can study over the long run: the cross-sectional sample and the supplemental black and Hispanic samples. We exclude the economically disadvantaged non-black/non-Hispanic supplemental sample as it was discontinued in 1990 and the military supplemental sample, most of which was discontinued in 1984.

In addition to these restrictions, we impose additional restrictions based on data availability, as discussed above, on the sample when analyzing earnings.

References


- ACEMOGLU, D., AND D. AUTOR (2011): “Skills, Tasks and Technologies: Implications for Employment and Earnings,” vol. 4 of *Handbook of Labor Economics*, chap. 12, pp. 1043–1171. Elsevier.
- AUTOR, D. H., AND D. DORN (2013): “The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market,” *American Economic Review*, 103(5), 1553–1597.
- AUTOR, D. H., D. DORN, G. H. HANSON, AND J. SONG (2014): “Trade Adjustment: Worker-Level Evidence,” *The Quarterly Journal of Economics*, 129(4), 1799–1860.
- BLS (1986): *Occupational Outlook Handbook*.
- (1994): *Occupational Outlook Handbook*.
- (1996): *Occupational Outlook Handbook*.
- (2004): *Occupational Outlook Handbook*.
- (2006): *Occupational Outlook Handbook*.
- BÖHM, M. J. (2013): “Has Job Polarization Squeezed the Middle Class? Evidence from the Allocation of Talents,” CEP Discussion Papers dp1215, Centre for Economic Performance, LSE.
- BUREAU OF LABOR STATISTICS (1986): *Occupational Outlook Handbook, 1986-87 Edition: Bulletin of the United States Bureau of Labor Statistics, No. 2250*. US Department of Labor.
- (1996): *Occupational Outlook Handbook, 1996-97*. Bernan Assoc.
- (2006): *Occupational Outlook Handbook, 2006-07 Edition*. US Department of Labor.
- (2018a): “NLSY79 Errata for 1979–2014 Data Release,” .
- (2018b): “NLSY79 Errata for 1979–2012 Data Release,” .
- (2018c): “NLSY79 Topical Guide to the Data: Employment,” .
- (2018d): *Occupational Outlook Handbook, 2018-2019 Edition*. Bernan Press.
- DAHL, G. B., AND L. LOCHNER (2012): “The Impact of Family Income on Child Achievement: Evidence from the Earned Income Tax Credit,” *American Economic Review*, 102(5), 1927–56.
- EMEK BASKER, L. F., AND S. KLIMEK (2015): “Customer-Labor Substitution: Evidence from Gasoline Stations,” Discussion Paper 15-45, Center for Economic Studies.
- FEDERAL RESERVE BANK OF MINNEAPOLIS (2018): “Consumer Price Index, 1913-,” .
- HESS, A. E. (2014): “America’s disappearing jobs,” .
- JACOBSON, L. S., R. J. LALONDE, AND D. G. SULLIVAN (1993): “Earnings Losses of Displaced Workers,” *The American Economic Review*, 83(4), 685–709.
- KAMBOUROV, G., AND I. MANOVSKII (2009): “Occupational specificity of human capital,” *International Economic Review*, 50(1), 63–115.
- PISSARIDES, C. A. (2000): *Equilibrium Unemployment Theory, 2nd Edition*. The MIT Press.
- ROY, A. D. (1951): “Some Thoughts on the Distribution of Earnings,” *Oxford Economic Papers*, 3(2), 135–146.

RUGGLES, S., J. T. ALEXANDER, K. GENADEK, R. GOEKEN, M. B. SCHROEDER, AND M. SOBEK (2010): "Integrated Public Use Microdata Series: Version 5.0 [Machine-readable database]," Discussion paper, Minneapolis: University of Minnesota, 2010.

UNITED STATES BUREAU OF THE CENSUS (1980): *1980 Census of Population: Alphabetic Index of Industries and Occupations*.


VENERI, C. M. (1997): "Evaluating the 1995 occupational employment projections," *Monthly Labor Review*, 120.

C Appendix figures and tables




Notes: Coefficients on the declining indicator along with their 95-percent confidence intervals (robust to clustering by 1985 3-digit occupation) are displayed, where the regressions vary the sample, controls, and outcome variables. Coefficients are scaled by the mean of the outcome variable in each estimation sample. 'Post' refers to cumulative earnings 1986-2013. 'Pre' refers to the sum of earnings 1975 & 1980 for the middle and old, and earnings in 1980 for the young. We dropped the 1956-1960 birth cohorts as they did not reach age 25 by 1980, and for a similar reason we did not use 1975 earnings data for the young. 'Individual controls' are those used in column (2) of Table 4, and 'all controls' are the ones from column (6) in that table.

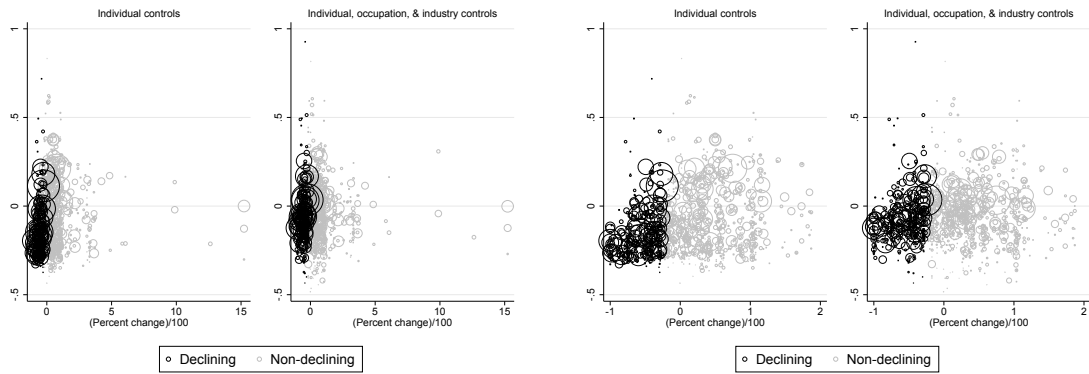
Figure OA2: Earnings prior to occupational decline



(a) Cumulative employment



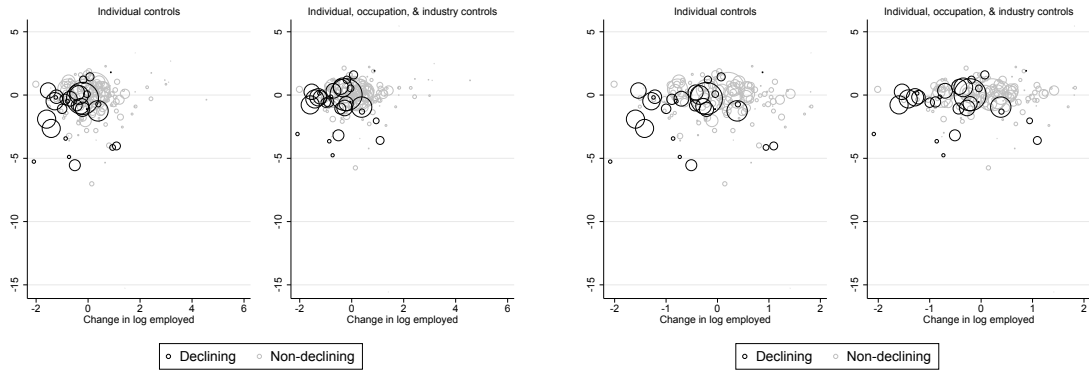
(b) Cumulative earnings



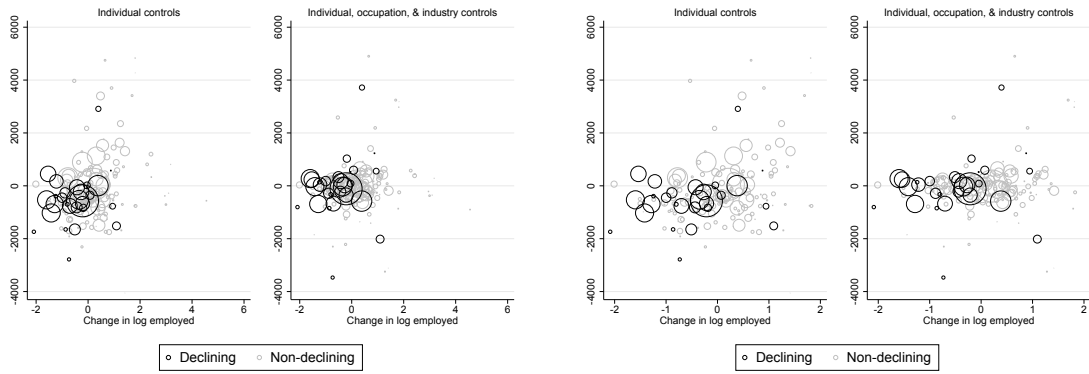
(c) Probability of remaining in the initial 3-digit occupation

Notes: Each bubble represents one of 1,052 5-digit Swedish occupations. Bubbles are scaled according to 1985 Swedish employment. The percent change in employment is assigned based on the changes 1984-2016 in the corresponding US occupation(s). Declining occupations are those that declined by more than 25 percent. Prior to aggregation, outcome variables were residualized based on the regression models in columns (2) and (6) in Tables 4 and 5, but with 'Declining' times its coefficient added (the mean difference between declining and non-declining occupations in the plots is thus exactly equal to the coefficients reported in the tables). The pairs of graphs on the right are truncated versions of those on the left.

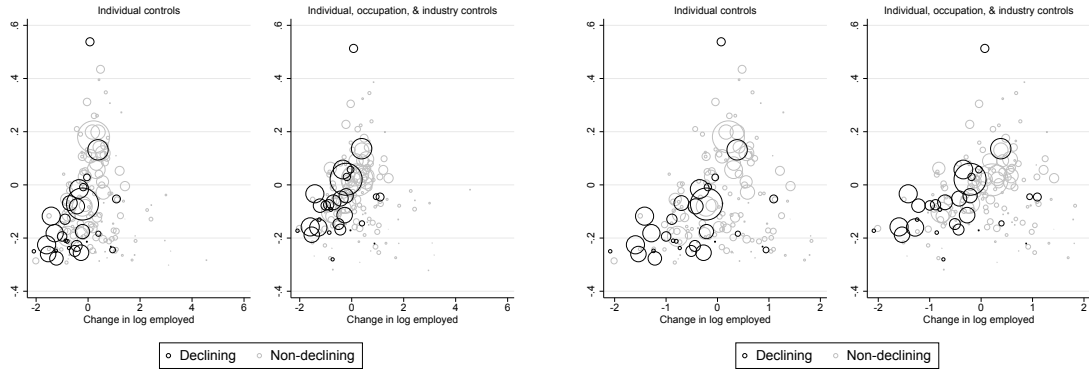
Figure OA3: Main outcomes and percent change in employment (US)



(a) Cumulative employment



(b) Cumulative earnings




(c) Probability of remaining in the initial 3-digit occupation

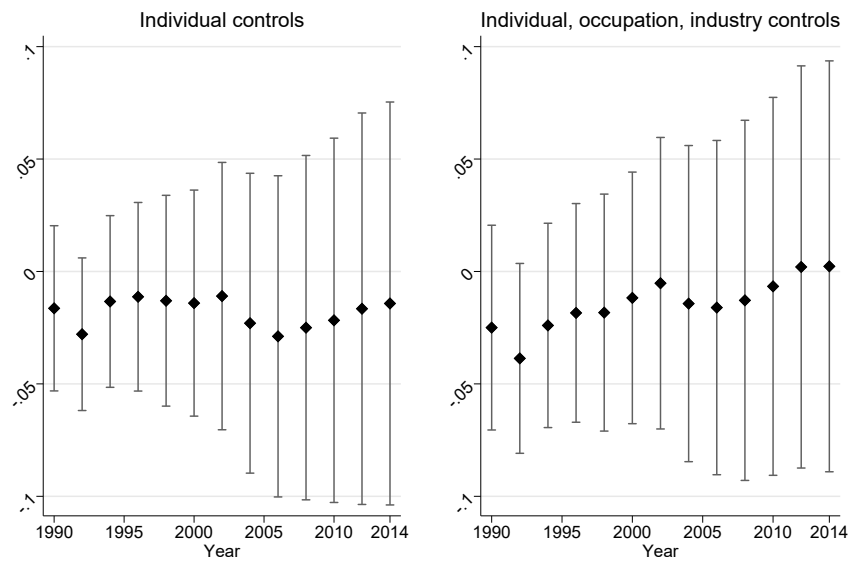
Notes: Each bubble represents one of 172 3-digit Swedish occupations. Bubbles are scaled according to 1985 Swedish employment. ‘Change in log employment’ refers to the actual change in log employment in each Swedish 3-digit occupation from 1985-2013. Occupations marked as declining are those in which more than two thirds of employment in 1985 was in a 5-digit occupation with the ‘Declining’ indicator equal to one. Prior to aggregation, outcome variables were residualized based on the corresponding regression models reported on in the last panel of Table OA3, with log employment change times its coefficient added (lines fitted to the plots would thus have slopes equal to the coefficients on log employment change reported in Table OA3). The pairs of graphs on the right are truncated versions of those on the left.

Figure OA4: Main outcomes and change in log employment (Sweden)

Cumulative earnings (2014 USD)



Cumulative earnings, divided by mean



Notes: Diamonds mark the coefficients on the declining indicator from the regression specifications reported in columns (2) and (6) of Table OA12, except that income accumulates only through time t . Capped bars indicate 95-percent confidence intervals.

Figure OA5: US (NLSY) occupational decline and individual-level earnings over time

Table OA1: Employment growth in Swedish 3-digit occupations 1985-2013

	(1)	(2)	(3)	(4)	(5)	(6)
Declining	-0.76 (0.17)				-0.44 (0.18)	-0.46 (0.18)
Employment share 1985		-1.23 (1.61)			-2.40 (1.57)	-2.31 (1.53)
Employment growth 1960-85		0.34 (0.08)			0.16 (0.09)	0.15 (0.08)
Predicted growth index			0.31 (0.07)		0.22 (0.08)	
Prediction: no change				-0.05 (0.44)		0.09 (0.42)
Prediction: increase, slow				0.46 (0.36)		0.25 (0.31)
Prediction: increase, average				0.74 (0.29)		0.55 (0.25)
Prediction: increase, fast				1.13 (0.29)		0.82 (0.28)
R^2	0.12	0.15	0.21	0.22	0.29	0.29

Notes: The dependent variable is the difference in log employment in Swedish 3-digit occupations between 2013 and 1985. 'Declining' is a binary variable at the level of 1985 Swedish 5-digit occupations indicating employment losses of more than 25 percent over the following three decades in the corresponding US occupation(s). The indicator has been collapsed to the 3-digit level and is thus a continuous regressor. The decline indicator and predictions have been constructed using the Occupational Outlook Handbook (various years). Regressions are weighted by 1985 Swedish employment shares. The number of observations is 172. Robust standard errors in parentheses.

Table OA2: Alternative cutoffs for occupational decline

	Employment		Earnings		Earnings, normalized		Remain	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Percent change $\in [-100, -50]$	-0.34 (0.20)	-0.18 (0.15)	-248.1 (115.6)	-90.0 (75.7)	-2.44 (0.62)	-0.98 (0.43)	-0.18 (0.040)	-0.10 (0.020)
Percent change $\in [-100, -25]$ (baseline)	-0.49 (0.20)	-0.19 (0.14)	-346.6 (120.3)	-126.4 (58.3)	-2.10 (0.53)	-1.11 (0.36)	-0.11 (0.041)	-0.045 (0.020)
Percent change $\in [-100, 0]$	-0.043 (0.20)	-0.0030 (0.13)	-35.0 (158.8)	-57.5 (74.7)	-0.70 (0.70)	-0.91 (0.47)	-0.15 (0.041)	-0.063 (0.021)
Percent change $\in [-100, 31]$ (below median)	0.14 (0.18)	0.15 (0.13)	-46.5 (150.7)	-61.9 (76.1)	-0.55 (0.57)	-0.53 (0.50)	-0.087 (0.037)	-0.0094 (0.022)
Baseline; control: percent change $\in (-25, 31)$	-0.72 (0.22)	-0.27 (0.16)	-460.5 (123.3)	-126.6 (61.9)	-2.40 (0.51)	-1.17 (0.40)	-0.077 (0.038)	-0.053 (0.018)
Individual controls	✓	✓	✓	✓	✓	✓	✓	✓
Occupation & industry controls		✓		✓		✓		✓
Observations				877,324				553,786

Notes: Results from regressions of various outcomes on indicators for occupational employment changes to lie in the indicated ranges are shown. Each panel represents a separate set of regressions. The underlying variable is the percentage change in employment for the US occupation(s) corresponding to the Swedish 5-digit occupation that the individual worked in during 1985. The last panel only keeps observations with a percentage change below the median, and the number of observations is thus halved. Normalized earnings are cumulative earnings divided by initial predicted earnings. See the notes to Tables 4 and 5 for further descriptions of variables and sample definitions. Robust standard errors, clustered by 1985 3-digit occupation, in parentheses.

Table OA3: Using continuous occupational employment changes as regressors

	Employment		Earnings		Earnings, normalized		Remain	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Percent employment change / 100 (US)	-0.019 (0.037)	-0.026 (0.036)	103.7 (30.2)	64.7 (14.9)	0.47 (0.11)	0.25 (0.13)	0.0058 (0.0068)	-0.0020 (0.0029)
Percent employment change / 100 (US), winsorized	0.010 (0.11)	0.000027 (0.080)	83.8 (112.0)	91.1 (47.5)	0.86 (0.40)	0.46 (0.25)	0.051 (0.025)	0.0035 (0.014)
Log employment change (SWE)	-0.034 (0.15)	0.049 (0.11)	306.4 (135.1)	73.7 (65.9)	0.85 (0.50)	0.087 (0.50)	0.11 (0.031)	0.066 (0.017)
Individual controls	✓	✓	✓	✓	✓	✓	✓	✓
Occupation & industry controls		✓		✓		✓		✓
Observations			877,324				553,786	

Notes: Results from regressions of various outcomes on change in occupational employment are shown. Each panel represents a separate set of regressions. 'Percent employment change (US)' refers to the percentage change in employment 1984-2016 for the US occupation(s) corresponding to the Swedish 5-digit occupation that the individual worked in during 1985. The winsorized measure of this variable top-codes changes at plus 217 percent (the 95th percentile). 'Log employment change (SWE)' refers to the change in log number employed 1985-2013 in the Swedish 3-digit occupation that the individual works in during 1985. Normalized earnings are cumulative earnings divided by initial predicted earnings. See the notes to Tables 4 and 5 for further descriptions of variables and sample definitions. Robust standard errors, clustered by 1985 3-digit occupation, in parentheses.

Table OA4: Occupational decline and individual-level cumulative employment and earnings 1986-2013—‘doughnut’ specifications

	(1)	(2)	(3)	(4)	(5)	(6)
<i>A. Cumulative years employed 1986-2013 (mean: 23.5)</i>						
Declining	-1.46 (0.53)	-0.97 (0.42)	-0.97 (0.42)	-0.82 (0.46)	-0.35 (0.28)	-0.41 (0.29)
<i>B. Cumulative real earnings ('000 2014 SEK) 1986-2013 (mean: 6,612)</i>						
Declining	-484 (608)	-403 (196)	-333 (177)	-140 (181)	-81 (158)	-217 (167)
<i>C. Cumulative real earnings divided by predicted initial earnings (mean: 39.2)</i>						
Declining	-5.40 (1.33)	-2.49 (1.09)	-2.56 (0.98)	-1.81 (1.07)	-1.18 (0.82)	-1.69 (1.05)
Demographics & earnings		✓	✓	✓	✓	✓
Life-cycle profiles			✓	✓	✓	✓
Predictors of growth				✓	✓	✓
Occupation dummies					✓	✓
Industry dummies						✓

Notes: Results from regressions of the indicated outcomes on a dummy for working in 1985 in a subsequently declining occupation are shown. The sample is the same as in Table 4, but excludes 3-digit occupations in which some but not all 5-digit occupations are coded as declining. Thus, within each 3-digit occupation, either all 5-digit sub-occupations decline, or none, leaving out intermediate cases (‘doughnut’). The number of observations is 488,484. Robust standard errors, clustered by 1985 3-digit occupation, in parentheses.

Table OA5: Occupational decline and individual occupational stability—‘doughnut’ specifications

	(1)	(2)	(3)	(4)	(5)	(6)
<i>A. Probability of working in same 3-digit occupation in 2013 as in 1985 (mean: 0.35)</i>						
Declining	-0.25 (0.046)	-0.21 (0.051)	-0.21 (0.052)	-0.12 (0.044)	-0.17 (0.046)	-0.10 (0.046)
<i>B. Probability of working in same 2-digit occupation in 2013 as in 1985 (mean: 0.40)</i>						
Declining	-0.21 (0.039)	-0.16 (0.045)	-0.16 (0.046)	-0.089 (0.043)	-0.12 (0.045)	-0.059 (0.042)
<i>C. Probability of working in same 1-digit occupation in 2013 as in 1985 (mean: 0.44)</i>						
Declining	-0.19 (0.036)	-0.14 (0.042)	-0.14 (0.043)	-0.077 (0.042)	-0.11 (0.043)	-0.045 (0.033)
Demographics & earnings		✓	✓	✓	✓	✓
Life-cycle profiles			✓	✓	✓	✓
Predictors of growth				✓	✓	✓
Occupation dummies					✓	✓
Industry dummies						✓

Notes: Results from regressions of the indicated outcomes on a dummy for working in 1985 in a subsequently declining occupation are shown. The sample is the same as in Table 5, but with the ‘doughnut’ restrictions from Table OA4 applied. The number of observations is 333,357. Robust standard errors, clustered by 1985 3-digit occupation, in parentheses.

Table OA6: Alternative functional forms for earnings

A. Discounted cumulative earnings												
	(1)	Discounted cumulative earnings		(4)	(5)	Discounted cumulative earnings, normalized						
		(2)	(3)			(6)	(8)					
Declining	-152.7 (57.1)	-47.8 (25.5)	-49.5 (24.4)	-33.2 (29.5)	-0.94 (0.25)	-0.47 (0.16)	-0.40 (0.18)					
Declining × rank			213.9 (68.5)			1.22 (0.25)						
Declining × bottom tercile				-166.5 (47.6)			-0.96 (0.23)					
Declining × top tercile				109.3 (62.0)			0.64 (0.21)					
Individual controls	✓	✓	✓	✓	✓	✓	✓					
Occupation & industry controls		✓	✓	✓		✓	✓					
Mean of dep. var.			3,476			19.4						
Mean of dep. var., bottom			2,954			17.5						
B. Rank, logs, and growth												
	Percentile rank in cumulative earnings		Logarithm of cumulative earnings		Percent growth in earnings 1985-2013							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Declining	-1.48 (0.84)	-0.85 (0.54)	-0.85 (0.50)	-0.95 (0.63)	-0.060 (0.022)	-0.021 (0.013)	-0.026 (0.014)	-0.00054 (0.017)	-41.8 (11.2)	-11.7 (8.61)	-9.39 (8.80)	0.73 (7.99)
Declining × rank			5.15 (0.93)				0.17 (0.035)				145.6 (35.7)	
Declining × bottom tercile				-3.26 (0.78)				-0.15 (0.037)				-110.0 (32.7)
Declining × top tercile				3.41 (0.89)				0.072 (0.018)				69.9 (24.9)
Individual controls	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Occupation & industry controls		✓	✓	✓		✓	✓	✓		✓	✓	✓
Mean of dep. var.			50.5				8.6				178	
Mean of dep. var., bottom			43.0				8.4				328	

Notes: Results from regressions of the indicated earnings measures on the declining indicator, within-occupation earnings rank or tercile dummies (coefficients omitted from table), and their interactions are shown. All regressions control for the level of 1985 earnings, with the exception of rank and logarithm as the outcome variables, in which case 1985 earnings rank and log of 1985 earnings are controlled for, respectively. Discounted cumulative earnings are calculated using an interest rate of 5 percent. Normalized earnings are cumulative earnings divided by initial predicted earnings. See the notes to Tables 4 and 5 for further descriptions of variables and sample definitions. The number of observations is 877,324, except when the log of cumulative earnings is the outcome variable, in which case the number is 875,830. Robust standard errors, clustered by 1985 3-digit occupation, in parentheses.

Table OA7: Heterogeneity by within-occupation residualized earnings rank

	Employment		Earnings		Earnings, normalized		Remain	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
A. Linear interaction								
Declining	-0.59 (0.22)	-0.20 (0.14)	-332.3 (90.2)	-154.0 (59.3)	-2.32 (0.56)	-1.22 (0.37)	-0.11 (0.041)	-0.042 (0.020)
Declining × rank	0.92 (0.33)	0.96 (0.29)	407.7 (141.9)	439.5 (137.3)	2.33 (0.59)	2.41 (0.56)	-0.020 (0.016)	-0.014 (0.015)
B. Dummy interactions								
Declining	-0.26 (0.22)	0.048 (0.16)	-302.5 (96.4)	-94.5 (62.9)	-1.94 (0.52)	-0.92 (0.38)	-0.095 (0.050)	-0.032 (0.025)
Declining × bottom tercile	-1.16 (0.36)	-1.11 (0.33)	-370.4 (93.5)	-390.7 (86.2)	-2.14 (0.48)	-2.10 (0.44)	-0.015 (0.019)	-0.0082 (0.017)
Declining × top tercile	0.16 (0.15)	0.24 (0.15)	220.2 (109.2)	202.7 (111.3)	0.99 (0.47)	1.12 (0.45)	-0.037 (0.027)	-0.026 (0.020)
Individual controls	✓	✓	✓	✓	✓	✓	✓	✓
Occupation & industry controls		✓		✓		✓		✓
Mean of dep. var.		23.4		6,926		38.7		0.29
Mean of dep. var., bottom		22.0		6,139		34.3		0.26
Observations				877,324				553,786

Notes: The notes to Table 6 apply, with the only difference that rank and terciles refer to the within-occupation distribution of 1985 earnings residualized by gender, cohort, and county. Robust standard errors, clustered by 1985 3-digit occupation, in parentheses..

Table OA8: Cumulative earnings of leavers and stayers in declining and non-declining occupations

	(1)	(2)	(3)	(4)
<i>A. All workers (553,169 observations)</i>				
Remain	335 (122)	303 (91)	305 (133)	284 (101)
Declining			-272 (122)	-127 (90)
Declining \times remain			177 (239)	190 (185)
<i>B. Employed in 2013 (404,043 observations)</i>				
Remain	-398 (115)	-498 (66)	-439 (124)	-531 (72)
Declining			-357 (123)	-188 (94)
Declining \times remain			238 (231)	312 (158)
<i>C. Employed in 2013, bottom third (140,892 observations)</i>				
Remain	-109 (139)	-285 (85)	-133 (143)	-307 (80)
Declining			-418 (145)	-238 (173)
Declining \times remain			-32 (596)	235 (425)
Individual controls	✓	✓	✓	✓
Occupation & industry controls		✓		✓

Notes: The dependent variable is cumulative earnings 1986-2013 in thousands of 2014 SEK. 'Remain' is an indicator for working in the same 3-digit occupation in 2013 as in 1985. Individual-level controls include female, cohort, county, and education dummies, as well as earnings in 1985. Occupation and industry controls include predicted life-time income, predictors of occupational growth, occupation dummies, and industry dummies. The sample is the same as that in Table 5, except for the restrictions indicated. Sampling weights are applied. Robust standard errors, clustered by 1985 3-digit occupation, in parentheses.

Table OA9: Baseline characteristics of workers in subsequently declining occupations—technology-related declines

	(1) Female	(2) Age	(3) Compulsory school	(4) High school	(5) College	(6) Earnings	(7) Manufacturing
<i>A. Occupational decline, pooled</i>							
Intercept	0.52 (0.078)	39.5 (0.41)	0.33 (0.030)	0.56 (0.033)	0.11 (0.027)	191.3 (10.8)	0.25 (0.050)
Declining	-0.25 (0.088)	-0.89 (0.63)	0.13 (0.035)	-0.063 (0.034)	-0.070 (0.028)	-0.23 (11.0)	0.38 (0.085)
<i>B. Occupational decline, by presence of technology link</i>							
Intercept	0.52 (0.078)	39.5 (0.41)	0.33 (0.030)	0.56 (0.033)	0.11 (0.027)	191.3 (10.8)	0.25 (0.050)
Declining	-0.32 (0.10)	0.033 (0.87)	0.13 (0.056)	-0.086 (0.051)	-0.041 (0.035)	5.31 (15.0)	0.26 (0.10)
Declining (technology)	0.11 (0.097)	-1.49 (1.01)	0.010 (0.059)	0.037 (0.050)	-0.047 (0.025)	-8.90 (14.6)	0.20 (0.12)

Notes: Results from OLS regressions of various baseline (1985) characteristics on a constant and indicators for working in a declining occupation are shown (see the notes to Table OA11 for a description of these indicators). Earnings are measured in thousand Swedish crowns inflated to 2014 levels. The sample is the same as in panel A of Table 2. The number of observations is 3,061,051. Robust standard errors, clustered by 1985 3-digit occupation, in parentheses.

Table OA10: Occupational decline and individual-level cumulative employment and earnings 1986-2013—technology-related decline

	(1)	(2)	(3)	(4)	(5)
<i>A. Change in log employment 1985-2013 in 3-digit occupation</i>					
Declining	-0.43 (0.16)	-0.25 (0.09)			
Declining (technology)	0.06 (0.17)	0.05 (0.13)	-0.38 (0.12)	-0.21 (0.14)	
<i>B. Cumulative years employed 1986-2013 (mean: 23.4)</i>					
Declining	-0.93 (0.44)	-0.45 (0.24)			
Declining (technology)	0.72 (0.45)	0.42 (0.23)	-0.21 (0.16)	0.01 (0.14)	-0.16 (0.12)
Declining (tech) × rank					1.31 (0.36)
<i>C. Cumulative real earnings ('000 2014 SEK) 1986-2013 (mean: 6,926)</i>					
Declining	-426 (232)	-181 (93)			
Declining (technology)	128 (262)	87 (102)	-303 (131)	-107 (65)	-122 (61)
Declining (tech) × rank					491 (155)
<i>D. Probability of working in same 3-digit occupation in 2013 as in 1985 (mean: 0.29)</i>					
Declining	-0.077 (0.051)	-0.029 (0.022)			
Declining (technology)	-0.058 (0.044)	-0.025 (0.029)	-0.135 (0.043)	-0.053 (0.026)	-0.056 (0.026)
Declining (tech) × rank					0.019 (0.016)
Individual controls	✓	✓	✓	✓	✓
Occupation & industry controls		✓		✓	✓
Observations (population—sample)	877,324—553,786		836,057—532,421		

Notes: Results from regressions of the indicated outcomes on indicators for working in 1985 in a subsequently declining occupation are shown (see the notes to Table OA11 for a description of these indicators). Columns (1)–(2) are based on the same samples as the results in Tables 4 and 5. Columns (3)–(5) exclude workers in occupations that are classified as declining without a technology link. Individual-level controls include female, cohort, county, and education dummies, as well as earnings in 1985. Occupation and industry controls include predicted life-time income, predictors of occupational growth, occupation dummies, and industry dummies. Sampling weights are used in the regression reported in panel C. Robust standard errors, clustered by 1985 3-digit occupation, in parentheses.

Table OA11: Employment growth in Swedish 3-digit occupations 1985-2013—technology-related declines

	(1)	(2)	(3)	(4)	(5)	(6)
Declining	-0.76 (0.17)	-0.44 (0.18)	-0.92 (0.27)	-0.37 (0.27)		
Declining (technology)			0.27 (0.33)	-0.11 (0.35)	-0.69 (0.20)	-0.49 (0.25)
Employment share 1985		-2.40 (1.57)		-2.41 (1.57)		-2.28 (1.61)
Employment growth 1960-85		0.16 (0.09)		0.16 (0.09)		0.16 (0.09)
Predicted growth index		0.22 (0.08)		0.23 (0.09)		0.22 (0.09)
R^2	0.12 172	0.29 172	0.12 172	0.29 172	0.06 148	0.22 148
Observations						

Notes: The dependent variable is the difference in log employment in Swedish 3-digit occupations between 2013 and 1985. 'Declining' is a binary variable at the level of 1985 Swedish 5-digit occupations indicating employment losses of more than 25 percent over the following three decades in the corresponding US occupation(s). 'Declining (technology)' indicates that this decline is related to technological replacement. Both indicators have been collapsed to the 3-digit level and are thus continuous regressors. Columns (10) and (11) exclude 3-digit occupations where 'Declining' is larger than or equal to 0.5 and 'Declining (technology)' is smaller than 0.5. Decline indicators and predictions have been constructed using the Occupational Outlook Handbook (various years). Regressions are weighted by 1985 Swedish employment shares. Robust standard errors in parentheses.

Table OA12: US (NLSY) occupational decline and individual-level earnings 1988-2013

	(1)	(2)	(3)	(4)	(5)	(6)
<i>A. Average real earnings (2014 USD) 1988-2013, no interpolation (mean: 44,083)</i>						
Declining	-2,661 (3,372)	-151 (1,589)	279 (1,635)	595 (1,750)	655 (1,584)	-24 (1,536)
<i>B. Average real earnings (2014 USD) 1989-2013, no interpolation, odd years only (mean: 46,057)</i>						
Declining	-2,600 (3,649)	123 (1,901)	607 (1,967)	954 (2,107)	1,124 (1,891)	384 (1,823)
<i>C. Average real earnings (2014 USD) 1988-2013 (mean: 46,891)</i>						
Declining	-2,970 (3,783)	-92 (2,029)	408 (2,100)	892 (2,255)	963 (2,012)	227 (1,969)
<i>D. Cumulative real earnings (2014 USD) 1988-2013 (mean: 1,216,117)</i>						
Declining	-95,964 (102,583)	-17,313 (55,596)	-5,801 (56,336)	12,310 (59,240)	23,322 (54,961)	2,783 (56,695)
<i>E. Cumulative real earnings divided by predicted initial earnings (mean: 44.2)</i>						
Declining	-4.04 (2.57)	-2.71 (2.44)	-3.69 (2.10)	-4.87 (3.45)	-2.10 (2.80)	-2.43 (2.52)
Demographics & earnings		✓	✓	✓	✓	✓
Life-cycle profiles			✓	✓	✓	✓
Predictors of growth				✓	✓	✓
Occupation dummies					✓	✓
Industry dummies						✓

Notes: Results from regressions of the indicated outcomes on a dummy for working in 1987 in a subsequently declining occupation are shown. Detailed descriptions of all variables and their construction are in the appendix; here, we summarize the main characteristics. Demographic controls include female, region, education and birth year dummies, and 'earnings' refers to the level of labor income in 1987. Occupation-level life-cycle profiles are cumulative earnings calculated for each individual based on their 1987 occupation. Predictor of growth is the 1986 OOH outlook for each individual based on their 1987 occupation. Occupation and industry dummies are at the broad group and group category levels, respectively. The sample includes all individuals with an occupation listed 1987 and at least 8 years of reported labor earnings. Sampling weights are applied. The number of observations is 6,679 in panels A-C and 5,817 in panels D and E. Robust standard errors, clustered by 1987 occupation, in parentheses.

Table OA13: US (NLSY) occupational decline and individual employment

	(1)	(2)	(3)	(4)	(5)	(6)
<i>A. Fraction of reported weeks in employment status (mean: 0.83)</i>						
Declining	0.014 (0.012)	0.017 (0.0091)	0.018 (0.0090)	0.014 (0.0091)	0.0035 (0.0089)	0.0042 (0.0094)
<i>B. Fraction of reported weeks in unemployment status (mean: 0.03)</i>						
Declining	0.0063 (0.0038)	0.0043 (0.0030)	0.0041 (0.0030)	0.0052 (0.0030)	0.0051 (0.0028)	0.0042 (0.0030)
Demographics & earnings		✓	✓	✓	✓	✓
Life-cycle profiles			✓	✓	✓	✓
Predictors of growth				✓	✓	✓
Occupation dummies					✓	✓
Industry dummies						✓

Notes: Results from regressions of the indicated outcomes on a dummy for working in 1987 in a subsequently declining occupation are shown. Detailed descriptions of all variables and their construction are in the appendix; here, we summarize the main characteristics. Demographic controls include female, region, education and birth year dummies, and 'earnings' refers to the level of labor income in 1987. Occupation-level life-cycle profiles are cumulative earnings calculated for each individual based on their 1987 occupation. Predictor of growth is the 1986 OOH outlook for each individual based on their 1987 occupation. Occupation and industry dummies are at the broad group and group category levels, respectively. The sample includes all individuals with an occupation listed 1987 and at least 418 weeks (8 years) of reported labor force status. Sampling weights are applied. The number of observations is 6,722. Robust standard errors, clustered by 1987 occupation, in parentheses.

Table OA14: US (NLSY) occupational decline and individual occupational stability

	(1)	(2)	(3)	(4)	(5)	(6)
<i>A. Probability of working in same occupation in 2014 as in 1987 (mean: 0.09)</i>						
Declining	-0.043 (0.017)	-0.039 (0.017)	-0.038 (0.017)	-0.032 (0.017)	-0.022 (0.017)	-0.012 (0.019)
<i>B. Probability of working in same occupation group in 2014 as in 1987 (mean: 0.20)</i>						
Declining	-0.0050 (0.032)	0.00079 (0.028)	0.0025 (0.029)	-0.0087 (0.031)	0.016 (0.029)	0.022 (0.026)
<i>C. Probability of working in same broad occupation group in 2014 as in 1987 (mean: 0.36)</i>						
Declining	-0.066 (0.052)	-0.048 (0.042)	-0.041 (0.043)	-0.044 (0.044)	0.033 (0.027)	0.041 (0.026)
Demographics & earnings		✓	✓	✓	✓	✓
Life-cycle profiles			✓	✓	✓	✓
Predictors of growth				✓	✓	✓
Occupation dummies					✓	✓
Industry dummies					✓	✓

Notes: Results from regressions of the indicated outcomes on a dummy for working in 1987 in a subsequently declining occupation are shown. Detailed descriptions of all variables and their construction are in the appendix; here, we summarize the main characteristics. Demographic controls include female, region, education and birth year dummies, and 'earnings' refers to the level of labor income in 1987. Occupation-level life-cycle profiles are cumulative earnings calculated for each individual based on their 1987 occupation. Predictor of growth is the 1986 OOH outlook for each individual based on their 1987 occupation. Occupation and industry dummies are at the broad group and group category levels, respectively. The sample includes all individuals with an occupation listed 1987 who were interviewed in 2014. Sampling weights are applied. The number of observations is 5,749. Robust standard errors, clustered by 1987 occupation, in parentheses.