Individual Consequences of Occupational Decline*

Per-Anders Edin Tiernan Evans Georg Graetz Sofia Hernnäs

Guy Michaels

November 24, 2018

Abstract

The prospect of labor-replacing technologies raises concerns about earnings and employment losses that workers may suffer when demand for their occupations declines. We estimate these losses using a new methodology that measures unanticipated declines in occupational employment, which we apply to panel data on individual workers in Sweden. When we compare workers with very similar initial characteristics, we find that on average 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 earnings losses. These earnings losses are partly accounted for by reduced time spent in employment, and increased time in unemployment and retraining.

KEYWORDS: Technological change, Occupations, Inequality

JEL CLASSIFICATION: O33, J24, J62

*Corresponding 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 the AEA meetings, CEP Policy Roundtable, Employment in Europe conference in Cyprus, Essex Workshop on Innovation and Labor Markets, ENEF conference in Sussex, Nordic Summer Institute in Labour Economics, Queen Mary Economic Workshop, and Royal Society and British Academy workshop; and seminar participants at the Institute of Industrial Economics in Stockholm, Konstanz, LSE, MIT, the Sveriges Riksbank, University of British Columbia, 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. Part of the work on this paper was done while Graetz visited ifo Institute Munich, and he thanks its faculty and staff for their hospitality. Any opinions and errors are our own.

1 Introduction

What are the long run employment and earnings losses for 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, 2018), 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 jeopardize 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. So it is 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).

Previous research has studied how workers cope with various adverse shocks. For example, a well-established literature studies the individual consequences of mass layoffs (for instance Jacobson, LaLonde, and Sullivan, 1993) and trade shocks (for instance Autor, Dorn, Hanson, and Song, 2014). But measuring technology's impact on individuals has proved harder. Following the seminal work of Autor, Levy, and Murnane (2003), much of the literature has focused on the task content of occupations, and especially on the distinction between routine and non-routine work. In this vein, Cortes (2016) and Autor and Dorn (2009) study panel data on individuals who differ in the extent of routine work that they do. We contribute to this literature by studying demand shocks across detailed occupations rather than broad tasks. We show that even similar occupations, such as typists and secretaries, may experience very different employment changes.[1] This variation lets us examine the consequences of occupational decline for workers with similar initial characteristics and in similar occupations.

To frame our empirical analysis of the consequences of occupational decline, we begin with the benchmark for studying sorting by comparative advantage—the Roy model (Roy, 1951). Our baseline frictionless model predicts that the earnings losses from occupational decline and the probability of remaining in a declining occupation both increase in initial earnings, and that the earnings losses are larger for those who remain. As we discuss below, these predictions are inconsistent with our findings, and

[1] Employment of typists has nearly vanished, while that of secretaries continues to grow.

we show how these results can be reconciled by including heterogeneous moving costs and involuntary displacement in the model.

After presenting our theoretical framework, we develop a new methodology for measuring 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 there were probable technology drivers to the declines we find; and to separate declines that were anticipated at the time from ones that were not. We then match this occupation-level information to individual-level, longitudinal data on the entire Swedish population. The granularity of occupations in the OOH and in Sweden's 1985 occupational classification allows us to separate similar occupations that experienced different employment changes.[2] Thus we utilize the best aspects of both countries' data: the US data allow us to better study occupational changes over time and separate anticipated and unanticipated changes, while the Swedish data allow us to follow a large number of individuals over time who differ in their exposure to occupational declines, depending on their occupations in 1985.

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, prior income, and location. In some specifications we further control for occupation-varying life-cycle profiles, predictors of occupational change, and fixed effects for broad occupations and industries.

We confirm that our OOH-based measure of occupational decline and the predicted changes in US employment both correlate with the employment changes in Sweden. Specifically, we find that occupations that we classify as declining in the US also declined in Sweden; compared to other Swedish occupations, their employment change was around 75 log points lower.[3]

We find that compared to workers with similar individual characteristics, those exposed to occupational decline lost about 5 percent (2 percent) of mean cumulative pre-tax earnings (employment). But compared to similar workers in similar occupations and industries, the cumulative earnings (employment) losses were only around 2 percent (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

[2]Sweden's occupational classification changes over time, so we lose much of the granularity when we study Swedish occupational employment growth from the mid-1980s to today.

[3]While the estimate is large and statistically significant, it is not precise enough for us to use the US declines to instrument Swedish decline, so we focus on reporting reduced-form estimates, as we further explain below.

declining occupations.

While mean earnings losses from occupational decline were limited, those in the bottom tercile of their occupation's earnings distribution in 1985 suffered larger losses, amounting to 8-11 percent of their mean earnings. 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. These findings show that the distributional consequences of occupational decline were economically meaningful, even if mean losses were limited.

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, our findings suggest that occupational decline led to slightly earlier retirement among middle-aged (in 1985) workers.[4]

We show that our main results on employment and earnings losses are robust to a battery of checks, including the use of different functional forms and thresholds for our key variables. We provide further evidence on the validity of our identification by showing that conditional on our set of controls, the workers in declining occupations had similar earnings both in the short run (from 1985-1990) and before 1985.

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 find similarly moderate estimates of mean earnings and employment losses (or in some cases even no losses) from occupational decline when we use micro data from the US (National Longitudinal Survey of Youth 1979), although the estimates, at least for earnings, are less precise than in the Swedish micro data. This suggests that the fairly moderate mean losses from occupational decline may generalize to settings beyond Sweden.

Our empirical findings inform our theoretical understanding of the responses to occupational decline in several ways. First, our model assumes that occupational declines are demand driven, and our findings that they involve both relative earnings losses and more occupational exits are consistent with this assumption. Second, our findings show that the largest earnings losses from occupational decline in Sweden are incurred by those who earned the least within their starting occupations. This finding is inconsistent with the frictionless version of the model, but it is consistent with the version where occupational switching costs decline in the workers' ability in the destination occupation. Moreover, our

[4]In companion work we study how occupational decline matters for other socio-economic outcomes, including health, family composition, geographic location, and welfare transfers.

empirical analysis sheds light on the nature of the occupational switching costs: 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.

Third, involuntary displacement may also play an important role. Our model can account for several of the empirical findings when we allow for involuntary displacement, in addition to heterogeneous switching costs. 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 switching probabilities are U-shaped in initial earnings, whereby low-earning workers switch occupations if displaced, while high-earning workers switch regardless of displacement.

Our paper is related to the literature on the extent of displacement that we may expect from technological change over the next few decades. These forecasts of occupational displacement range from almost 50 percent (Frey and Osborne, 2017) to around 10 percent (Arntz, Gregory, and Zierahn, 2016). Others conclude that computers (Bessen, 2016) and automation more generally (Autor and Salomons, 2018) have, at least so far, not been responsible for net employment losses. Despite these differences, however, economists generally agree that some jobs will be replaced by technology, and our study offers a methodology for quantifying the losses from occupational displacement.

Our paper is also related to contemporaneous and independent work (Schmillen, 2018), who studies employment outcomes for German trainees who are hit by occupational demand shocks. While our studies differ in their research questions, the level of occupational variation, the econometric inference, and the outcomes considered, we both find modest employment losses from occupational shocks. Our paper is also related to independent work by Galaasen and Kostøl (2018), who find that short run losses of displaced workers are related to occupation-specific demand.

When we compare the magnitudes of earnings losses across studies we should, of course, be careful; there are important differences in the settings, the sample and restrictions used, the set of controls, and the duration for which people are followed up. With those caveats in mind, we compare our findings to a number of studies on the costs of job displacement. For example, studies 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 these losses 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). The lower mean losses that we find from occupational decline may be the result of a combination of factors: occupational decline is often gradual and some of it may be managed through retirement and less hiring; its gradual nature may allow some workers time to adjust through job-to-job moves without losing employment; and occupational decline need not be associated with significant adverse spillovers to the local economy, unlike mass layoffs (Gathmann, Helm, and Schönberg, 2018).[5]

Our paper is also related to the literature on the costs of moving across occupations. For example, Cortes and Gallipoli (2017), Kambourov and Manovskii (2009), Pavan (2011), and Sullivan (2010) study the human capital losses associated with switching occupations.[6] Our contribution to this literature is that we study occupational movement that is driven by plausibly exogenous demand shocks, and investigate how these costs vary by workers' initial position in their occupations' earnings distributions.[7]

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). 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), which may have contributed to the modest losses from occupational decline that we find.[8] At the same time, our finding of similarly modest mean earnings and employment losses from occupational decline in the US, suggests that our findings may generalize across different institutional settings.

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

[5] In the trade literature, Autor, Dorn, Hanson, and Song (2014) find that a worker in the third quartile of exposure to Chinese imports loses an average of 2.5 percent of initial earnings per year, compared to a similar worker in the first quartile. This figure is more similar to the losses that we find from occupational decline, and like us they also find that lower earning individuals suffer larger losses.

[6] 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.

[7] 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.

[8] Another 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.

2.1 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 every 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_{ki} 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.[9] 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 $\widetilde{\pi}_t \equiv \pi_{Bt} - \pi_{At}$ and assume for simplicity that $\widetilde{\pi}_1 = 0$.[10] Prices are determined in equilibrium by supply and demand. However, here we take them as given, and analyse 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.

In period 2, there is a negative demand shock to occupation A such that $\pi_{A2} - \pi_{A1} = -d$ and $\widetilde{\pi}_2 = d, d > 0$. This may be due to labor-replacing technology becoming available, or cheaper, in occupation A. Note that we assume d to be the equilibrium price change due to the shock. If humans and machines are perfect substitutes in occupation A, then a fall in the rental price of machines is fully passed through to occupational prices (see for instance Autor, Levy, and Murnane, 2003). If they are imperfect substitutes,

[9] 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.

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

however, and if labor supply to occupation A is upward sloping, then occupational prices will fall by less than the fall in the rental price of machines (see for instance Autor and Dorn, 2013). Thus, with imperfect substitutability, an upward sloping labor supply curve, as results here from occupational choice, already represents one mechanism that mitigates negative occupational demand shocks.

We are interested in the consequences of the shock for the earnings of workers who start out in occupation A. In particular, we define the earnings loss as earnings in the absence of the shock minus earnings after the shock occurred. The expected value of this loss (in logs), conditional on ability in the initial occupation, is a weighted average of the losses of those who switch occupations and those who remain, with the weights equalling the probabilities of each event,[11]

$$\mathbb{E}[\text{loss}|\alpha_{iA}] = \mathbb{P}(\text{switch}|\alpha_{iA})\mathbb{E}[\text{loss}^{\text{switch}}|\alpha_{iA}] + (1-\mathbb{P}(\text{switch}|\alpha_{iA}))\mathbb{E}[\text{loss}^{\text{remain}}|\alpha_{iA}]. \tag{1}$$

If there is no displacement and all moves are voluntary, then earnings losses of workers who remain are equal to the amount by which prices fall, $\mathbb{E}[\text{loss}^{\text{remain}}|\alpha_{iA}] = d$. Moreover, losses of those who move must be less than d by revealed preference, so that $\mathbb{E}[\text{loss}|\alpha_{iA}] < d$. Thus, earnings losses are mitigated by voluntary occupational switching.

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. Equation (1) shows that this depends on how both the probability of switching and the loss conditional on switching, vary with α_{iA}. In particular, for earnings *losses* to be *decreasing* in initial earnings it is sufficient that switching probabilities are *increasing*, earnings losses of switchers are *non-increasing*, and earnings losses of stayers are *non-decreasing*, in α_{iA}. 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 $\overline{\alpha}$. We argue below that our results are robust to alternative distributional assumptions.

2.2 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 1 (blue

[11] Equation (1) is conditioned on working in occupation A in period 1, but to avoid clutter this is not made explicit.

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 1. As the figure suggests, under independently and uniformly distributed skills, the fraction who switch among those initially working in A is weakly decreasing in α_{iA}. Moreover, we can show that switchers' earnings losses are also weakly increasing in α_{iA}; and that taken together, mean earnings losses for workers starting out in occupation A are strictly increasing in α_{iA}, and hence initial occupational earnings.[12]

An intuitive explanation for this result goes as follows. For the sake of the argument, call occupation A "typist" and occupation B "cashier", where typists are the ones who suffer a negative demand shock. The worst typists will 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.[13]

2.3 Costs of switching between occupations

Here we continue to assume that the period-2 price change is unanticipated, but now we suppose that there are costs of switching occupations. We think of these costs as the time lost searching for a new job, retraining, and perhaps moving geographical location, 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.

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. 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, $\mathbb{E}[\widetilde{\pi}_2] = \widetilde{\pi}_1 = 0$, where the last equality is due to our earlier simplifying assumption.[14] In this case, worker i chooses occupation A in period 1 if and only if $\alpha_{iA} \geq \alpha_{iB}$, as before. The workers who switch to occupation B after the price change satisfy the

[12] All formal derivations are presented in the appendix.

[13] A sufficient condition for earnings losses to be higher for the most able than for the least able is that there is positive probability mass everywhere in the $\alpha_A - \alpha_B-$plane and that support is finite.

[14] Instead of the random walk assumption we could impose that demand changes are somehow otherwise perfectly unforeseen, for instance due to adaptive expectations (in Section 2.5 we consider the case where demand changes are anticipated).

inequalities $\alpha_{iB} \leq \alpha_{iA}$ and $\alpha_{iB} > \alpha_{iA} - (d - c)$. Panel (b) of Figure 1 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). Thus, switching probabilities are decreasing and earnings losses are increasing in α_{iA}, as before. And as argued in Section 2.2 using the example of typists and cashiers, the conclusion regarding the "best" and "worst" in the declining occupation is robust to alternative distributional assumptions.

Suppose instead that workers who wish to switch from A to B must pay a switching cost equal to $C - \alpha_{iB}$, with $C > \overline{\alpha}$. (The condition $C > \overline{\alpha}$ 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.[15] If we continue to assume that occupational log prices follow a random walk, then we again obtain the result that worker i chooses occupation A in period 1 if and only if $\alpha_{iA} \geq \alpha_{iB}$. 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 1. 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. Furthermore, switchers' period-2 expected log earnings (taking into account switching costs in the form of lost time) are strictly increasing in period-1 earnings. Hence, mean losses conditional on initial earnings are (weakly) *decreasing* in initial earnings. In terms of the example above, 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 there will be many who have substantial earnings potential as cashiers, which in addition means that their switching costs are low. 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.

2.4 Job displacement

So far, we have been concerned with earnings losses as a function of initial earnings in the context of a simple 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

[15] In reality there are many more than just two occupations, and the best outside option for those at the top of declining occupations may be to move to a better job that they initially avoided (perhaps due to non-pecuniary factors); and this may less likely be the case for those at the bottom of declining occupations. This mechanism is outside the scope of our model, but its implications would arguably be similar to those of a heterogeneous switching cost.

of finding a new job in the initial occupation, may overturn this result.[16]

Suppose that workers who start in occupation A may experience job displacement, and incur a time cost $\widehat{c} > 0$ to find a job in A, or a cost c to find a job in B (we consider heterogeneous switching costs in the next paragraph). 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 $\alpha_{iB} > \alpha_{iA} - (d - (c - \widehat{c}))$, 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. 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.

In the case of a heterogeneous cost of moving to occupation B, $C - \alpha_{iB}$, we introduce in symmetric fashion a cost of finding a job in occupation A in case of displacement, $C - \alpha_{iA}$. 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. We again obtain the result that earnings losses of displaced, occupation-switching workers may be larger than those of stayers, by the same argument as above.

Another difference between the cases of constant and heterogeneous switching costs under displacement is how the probability of switching occupation, and the reason for switching, vary with initial earnings. For simplicity, let us assume that the risk of displacement is independent of initial earnings. In the case of a constant switching cost, the probability of a voluntary switch is decreasing, and that of a displacement-induced switch is hump-shaped, in initial earnings. The overall probability of switching is decreasing in initial earnings. In contrast, with heterogeneous switching costs the probability of a displacement-induced switch is decreasing, and that of a voluntary one is increasing in initial earnings. The overall probability of switching is U-shaped in initial earnings.[17]

[16] Recall 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).

[17] The last result is due to the assumption that the probability of displacement is less than one, so that not all workers in the yellow area switch. We also note that allowing for heterogeneous displacement probabilities, where lower ranked workers are more likely to be displaced, may help account for some of the patterns that we document.

2.5 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 $\widetilde{\pi}_2 = d$ 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 2.2. 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 in Section 2.3. Now we have that 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 the switching cost is large, $\beta c > d - c$, then all re-sorting in response to the anticipated shock occurs in period 1 already, before the shock hits, and no moves occur after period 1.

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.

2.6 Summary of theoretical results

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

Furthermore, our frictionless baseline model makes three strong 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 suggests that in our empirical work, we should not only focus 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. In the next sections we investigate this empirically by comparing lost time spent in employment to increased time spent unemployed, and by considering the impact of occupational decline on the retirement age.

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.

3 Data and empirical strategy

3.1 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 containing a sample of US residents, and from the 1980 US Census. Here we focus on the Swedish data and the OOH, and describe the remaining data sources in the appendix.

3.1.1 Swedish individual-level and occupation 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 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.[18] 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.

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.[19] 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).[20]

3.1.2 The Occupational Outlook Handbook, the mapping to Swedish occupations, and the definition of occupational decline

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

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

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

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

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.[21]

Using the reported employment numbers from our two editions of the OOH, we calculate the percentage growth in employment 1984-2016.[22] 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.[23] 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.[24]

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.[25]

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

[21] The 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.

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

[23] Between 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 appendix for details.

[24] Veneri (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.

[25] The details of the weighting scheme are given in the appendix.

using a number of alternative thresholds. We also use information from the OOH to determine whether technology likely played a role in the decline.[26] In 1985, 13 percent of Swedish employees worked in subsequently declining occupations, and 8 percent worked in subsequently declining occupations where the decline is likely linked to technology. We present comprehensive descriptive statistics for workers in declining and non-declining occupations in Section 3.1.4.

Having described our US-based data on occupational changes, the mapping to Swedish occupations, and the definition of occupational decline, we now investigate to what extent occupational changes and forecasts from the US correlate with the changes that took place in Sweden. 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. But we are able to directly assess the relevance that our US-based definition of occupational decline has for Swedish occupational employment growth. To do so, we collapse the declining indicator, as well as employment forecasts, to the 3-digit SSYK96 classification. Column (1) of Table 1 presents results from regressing Swedish occupational employment growth (log changes) 1985-2013 at the 3-digit level on the cell mean of the declining indicator, weighting the regression by 1985 Swedish employment shares. The difference in employment growth between 3-digit occupations which contain no declining sub-occupations and those in which all sub-occupations are classified as declining, is substantial at 76 log points.

We also investigate to what extent this decline was predictable in 1985. Column (2) shows that prior (1960-85) employment growth is a strong predictor of growth from 1985-2013—growth was persistent—while initial (1985) employment shares do not contribute additional explanatory power (the combined R-square is 0.15). Strikingly, the predicted growth index based on OOH forecast categories has even more explanatory power for Swedish employment growth (the R-square is 0.21, column (3)). Entering the OOH predictions as categorical variables only marginally improves the forecast, as seen in column

[26]To 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 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.

(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).

In sum, the results in Table 1 establish that during our sample period, occupational decline was correlated between the US and Sweden; that employment projections for US occupations were successful predictors of employment growth in the corresponding Swedish occupations; and that a substantial part of occupational decline could not be predicted in 1985, at least not by the variables at our disposal.

3.1.3 The rationale for our measure of occupational decline based on US data

Table 1 shows that it is possible to derive a measure of Swedish occupational decline from the (US) OOH employment changes. Here we explain why we believe it is 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 SSYK96 codes, and having more codes helps us get variation from small and declining occupations (such as typists, whose employment fell sharply, compared to secretaries, whose employment grew). To use the OOH data we need to match them to the YRKE5 codes, but since the YRKE5 are more numerous we do not lose much variation from this match. Second, since the SSYK96 codes were introduced from 1996 they reflect a judgement on the importance of occupations a decade after the start of the occupational decline we study. As such they are more likely to pool occupations with little employment by 1996 (including occupations that declined sharply by then) with non-declining occupations. In contrast, when we match OOH data over time we can separate occupations that have disappeared or almost disappeared, because the 2018-19 OOH contains detailed information on many smaller occupations. Finally, using occupational declines measured in Sweden on 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.

But this raises a further question: 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 the SSYK96 codes, we lose much of the variation on occupational decline because of the coarseness of the classifications and the lower likelihood of separating occupations in sharp decline. Second, using 2SLS exacerbates the problem, since we are using only part of the variation in the decline. Finally, as we discuss 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 affect 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 and non-parametric figures 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 some 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 reported in Table 1. 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.

3.1.4 Sample restrictions, construction of variables, and individual-level descriptive statistics

We next describe our 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.[27] 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

[27]There 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.

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.

Our main right-hand side variable—our treatment of interest—is an indicator for working in 1985 in a subsequently declining occupation, as defined above. 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 simply summing up years observed as employed and real annual labor earnings, obtaining the variables cumulative years employed and cumulative earnings.[28] 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.[29] 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 has left the labor force.[30] 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. 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.[31]

We end our data description by presenting 1985 individual-level descriptive statistics, thereby exploring the sorting of individuals into occupations that subsequently declined. Table 2 presents results

[28] We 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.

[29] The 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.

[30] Our 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.

[31] The 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.

from regressions of several individual characteristics on a constant and the declining indicator. The top panel of the table considers our starting sample, the working age population as a whole (ages 16-64) restricted to those individuals with non-missing demographic and labor market information, as described above. The bottom panel focuses on our baseline sample (ages 25-36). In both cases, the sorting patterns are similar: Those in occupations that subsequently declined were 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. There are no statistically or economically significant differences in age. In sum, workers starting out in subsequently declining occupations were systematically different from the rest of the workforce in terms of gender, education, and industry. As we discuss in the next subsection, we therefore have to control for these characteristics—and others—to get closer to a causal interpretation of our empirical results.

3.2 Empirical strategy

Our objective is to estimate the consequences of occupational decline for individual workers' careers. For concreteness, let D_i indicate that worker i starts out in an occupation that subsequently experiences an adverse shock and declines, and let y_i be some outcome measured after the shock has hit. In terms of the theoretical model from Section 2, $D_i = 1$ would indicate working in occupation A in the first period ($k_{i1} = A$), and y_i would refer to second-period log earnings.[32] Suppose we regress y_i on D_i. In potential outcomes notation (Rubin, 1974), let $y_i(1)$ denote the outcome of worker i if the negative shock has occurred, and let $y_i(0)$ denote the outcome for the same worker in the absence of the shock. The coefficient on the declining dummy in the probability limit equals

$$\begin{aligned}&\mathbb{E}\left[y_i|D_i=1\right]-\mathbb{E}\left[y_i|D_i=0\right]=\\&\qquad\underbrace{\mathbb{E}\left[y_i(1)|D_i=1\right]-\mathbb{E}\left[y_i(0)|D_i=1\right]}_{\text{LOSS}\times(-1)}+\underbrace{\mathbb{E}\left[y_i(0)|D_i=1\right]-\mathbb{E}\left[y_i(0)|D_i=0\right]}_{\text{SELECTION BIAS}}.\end{aligned}$$

It thus captures both the average treatment effect on the treated—for concreteness, the negative of the earnings loss for those in declining (relative to other) occupations, as defined in Section 2—and a selection bias term, which will be different from zero if workers starting out in the declining occupation would

[32] There are only two periods in the model, but as discussed, the second period may be interpreted as all periods following the shock. In the data, we cannot exactly pin down the timing of the shock. To be sure, we do not claim that it occurred precisely at the end of 1985.

have had different potential outcomes than those starting out in the non-declining occupation, had the shock not occurred. In our theoretical model, we assume the two occupations to be perfectly symmetric apart from the shock, so that $\mathbb{E}\left[y_i(0)|D_i = 1\right] = \mathbb{E}\left[y_i(0)|D_i = 0\right]$ (or, using the notation from the model, $\mathbb{E}\left[y_{i2}(0)|k_{i1} = A\right] = \mathbb{E}\left[y_{i2}(0)|k_{i1} = B\right]$).

In reality, of course, there are many reasons why the selection term may be non-zero, and we examine several of them in turn. First, as already discussed and documented in Table 2, workers in subsequently declining occupations in 1985 differed systematically from the rest of the workforce in terms of observable characteristics such as gender and education, and these may affect their outcomes irrespective of occupational decline. We therefore control for a rich set of observables, including income in 1985, in our regressions.

Second, declining occupations may differ in their life-cycle earnings profiles even in the absence of adverse shocks causing them to decline. To address this concern, we predict each worker's cumulative earnings 1986-2013 based on her initial 3-digit Swedish occupation, and use this prediction as a control.[33] Admittedly, the predictions are imperfect counterfactuals: many workers switch occupation at some point in their careers, but our prediction is based on the 1985 cross-section due to data limitations. Furthermore, there may be differences in labor market attachment across occupations, but our predictions necessarily are based on samples of employed workers. Nevertheless, an individual-level regression of actual cumulative earnings 1986-2013 on predicted earnings gives an R-square of 0.52 (the coefficient equals 0.98), indicating that predicted earnings absorb much variation, and are thus a useful control.

Third, it is possible that even conditional on the controls above, some individuals were better at anticipating the upcoming occupational and sorted away from declining occupations. To address this concern, we control for the 1986-87 OOH employment forecasts; as we saw above, even though these predictions concerned the US economy, they did well at forecasting Swedish occupational employment growth 1985-2013. In addition, we control for 1985 Swedish employment shares and 1960-85 Swedish employment growth at the level of 3-digit SSYK96 occupations, as these variables may have helped Swedish workers anticipate occupational decline. By controlling for all these factors we are likely iso-

[33] We calculate predicted cumulative earnings 1986-2013 as follows. We take all workers aged 16-64 in 1985 who earned at least one base amount. In each 3-digit SSYK96 occupation, we regress log earnings on a quartic in age and dummies for female, county, and education (we also run the same regression for all 2-digit SSYK96 occupations). We then predict the change in log earnings for all individuals and across all years 1986-2013 using the estimated quartic in age from a worker's 1985 3-digit occupation. If the 3-digit occupation contained fewer than 500 workers in 1985, we instead use the quartic estimated at the level of the corresponding 2-digit occupation to avoid measurement error due to imprecise estimates. We also add to workers' annual earnings growth the average wage growth in Sweden 1986-2013. We initialize each worker's log earnings using the prediction—involving the full set of covariates—from the occupation-level regressions. We then add to the initial level of earnings the annual earnings for each year 1986-2013 implied by the initial level and the predicted earnings growth (we give zero predicted earnings to workers older than 64).

lating the effects of surprise declines, which are unlikely to have induced sorting of Swedish workers in 1985.

Fourth, there may be unobserved differences between workers who sort across broad occupational groups (for instance, managers, professionals, or clerical workers) even conditional on all the variables mentioned above. To address this concern we control for 1-digit occupation dummies. Finally, industry-level shocks, possibly from changes in demand or increased trade openness, may have been correlated with occupational decline. To mitigate this concern we further add 2-digit industry dummies as controls. As with all other controls, these occupation and industry indicators are based on workers' jobs in 1985.

Adding the controls just discussed should get us closer to estimating the causal effect of occupational decline on workers' career outcomes, as they help us compare workers who are observationally similar across declining and non-declining occupations, as well as workers in declining occupations to those in observationally similar non-declining occupations. But introducing rich sets of controls also carries risks. Workers exiting the declining occupations may switch to a small set of otherwise similar non-declining occupations, where earnings even of incumbent workers may be depressed as a result of this inflow. This would lead us to underestimate the earnings effects of occupational decline. And even in the absence of such general equilibrium effects, employing a rich set of controls may lead our estimates to put more weight on groups of comparable occupations where there are roughly as many declining as non-declining sub-occupations. But losses may be lower in such cases, as workers have many non-declining occupations similar to the declining ones to choose from. In light of these concerns, we report results based on different sets of controls throughout.[34]

In light of the above discussion, we implement our empirical strategy by OLS estimation of the equation:

$$y_i = \beta_1 D_{k(i)} + \beta_2 \mathbf{x}_i + \beta_3 \widehat{e}_{i,k(i)} + \beta_4 \mathbf{w}_{k(i)} + \alpha_{s(i)} + \varepsilon_i. \tag{2}$$

The outcome y_i may be individual i's cumulative earnings, cumulative employment, or one of the other outcome variables described in Section 3.1.4. The variable $D_{k(i)}$ indicates whether individual i's 5-digit occupation k in 1985 subsequently declined, as defined in Section 3.1.2. Individual-level controls $\mathbf{x}_i$ include dummies for female, cohort, 1985 county of residence, seven education categories, as well as real income in 1985. $\widehat{e}_{i,k(i)}$ denotes predicted cumulative earnings 1986-2013 as described above. The vector

[34] In addition, we provide direct evidence on the importance of substitute occupations, by restricting our analysis to those 3-digit occupations in which either all or non of the sub-occupations declined.

of occupation-level controls $\mathbf{w}_{k(i)}$ includes predictors of occupational employment growth (1985 Swedish employment shares, 1960-85 Swedish occupational employment growth, and the predicted growth index from the OOH), and sometimes 1-digit occupation dummies.[35] $\alpha_{s(i)}$ is a fixed effect for the individuals' 1985 2-digit industrial sector. Finally, ε_i is an error term. While decline is measured at the level of (over 1,000) 5-digit occupations, we cluster standard errors at the level of (172) 3-digit occupations. We use this rather conservative approach because our measure of decline is derived from 401 US occupations in the 1986-7 OOH, so we cannot treat each 5-digit Swedish occupation as independent.

We also investigate whether the consequences of occupational decline vary by relative ability, as measured by workers' initial within-occupation earnings rank. We estimate:

$$y_i = \gamma_1 D_{k(i)} + \gamma_2 \text{rank}_{i,k(i)} + \gamma_3 D_{k(i)} \times \text{rank}_{i,k(i)} + \gamma_4 \mathbf{x}_i + \gamma_5 \widehat{e}_{i,k(i)} + \gamma_6 \mathbf{w}_{k(i)} + \alpha_{s(i)} + u_i, \tag{3}$$

where $\text{rank}_{i,k(i)}$ is individual i's rank in the distribution of 1985 earnings within her 3-digit SSYK96 occupation k. Alternatively, we include dummies for the bottom and top tercile as well as their interaction with the declining measures on the right-hand side; or we base the rank (and terciles) on the within-occupation distribution of residualized earnings if we are interested in a relative ability measure that is net of gender, life-cycle, and regional pay differences.[36]

The interaction coefficient γ_3 and corresponding coefficients on interactions with terciles, require different assumptions for a causal interpretation to be valid than the main coefficient γ_1. For simplicity, let us consider a dummy for being in the bottom tercile of initial within-occupation earnings, interacted with the declining dummy. The coefficient on this interaction in the probability limit equals

$$\mathbb{E}\left[y_i|D_i=1,T_1\right] - \mathbb{E}\left[y_i|D_i=1,T_{2,3}\right] - \left\{\mathbb{E}\left[y_i|D_i=0,T_1\right] - \mathbb{E}\left[y_i|D_i=0,T_{2,3}\right]\right\} =$$

$$\underbrace{\mathbb{E}\left[y_i(1)|D_i=1,T_1\right] - \mathbb{E}\left[y_i(0)|D_i=1,T_1\right]}_{\text{LOSS}\times(-1)\text{ for }T_1} - \underbrace{\left\{\mathbb{E}\left[y_i(1)|D_i=1,T_{2,3}\right] - \mathbb{E}\left[y_i(0)|D_i=1,T_{2,3}\right]\right\}}_{\text{LOSS}\times(-1)\text{ for }T_{2,3}}$$
$$+\underbrace{\mathbb{E}\left[y_i(0)|D_i=1,T_1\right] - \mathbb{E}\left[y_i(0)|D_i=1,T_{2,3}\right] - \left\{\mathbb{E}\left[y_i(0)|D_i=0,T_1\right] - \mathbb{E}\left[y_i(0)|D_i=0,T_{2,3}\right]\right\}}_{\text{SELECTION BIAS}},$$

[35]Recall from Section 3.1 that decline and the OOH predictions vary at the level of 5-digit (YRKE5) occupations, and that 1985 Swedish employment shares and 1960-1985 Swedish employment growth vary at the level of 2-digit (SSYK96) occupations. To avoid clutter, we use a single symbol k to denote occupation, even though the underlying classifications may differ.

[36]We include main effects in all our specifications, though we do not report them to avoid clutter, and also because they are difficult to interpret when individual earnings are controlled for also. In any case, the coefficients on the interaction terms do not change noticeably when excluding main effects, and the main effects are of small magnitude and statistically insignificant when including the full set of individual, occupation, and industry controls.

where T_1 ($T_{2,3}$) indicates being among the bottom earnings tercile (the top two) in the initial occupation. Considering again earnings as an example, the first line after the equality sign is the difference in the (negative of the) counterfactual earnings losses between bottom-tercile and the remaining workers who started out in the declining occupation. The interaction coefficient will estimate this difference in average treatment effects in the absence of selection bias. The assumption required for the selection bias term to disappear is similar to the parallel trends assumption in a differences-in-differences setting: in the absence of the shock, average differences in potential outcomes between the bottom tercile and the upper two terciles must be the same across the declining and non-declining occupations. This assumption may fail in practice: for instance, earnings dispersion may differ across occupations, as emphasized by Roy (1951). However, we will see that our estimates of interaction coefficients are robust to varying the set of controls.

4 Empirical analysis

In this section we present the findings from our empirical analysis. We first examine how career employment, earnings, and occupational mobility differed between workers who started out in a subsequently declining occupation and those who did not. Second, we investigate how the outcome of occupational decline differed by workers' initial within-occupation earnings rank. Third, we explore some of the channels through which occupational decline may have operated, including unemployment, retraining, and—for older workers—early retirement. Fourth, we examine what the outcomes for those who experienced occupational decline might have been in the absence of the decline. Fifth, we examine whether occupational declines with observed links to technology differed from the other declines in their implications. Sixth, to alleviate concerns that the consequences of occupational decline that we document are specific to Sweden, we repeat the main elements of our analysis using NLSY data from the US. At the end of this section, we offer an interpretation of our results through the lens of the theoretical model from Section 2.

4.1 Main results on employment, earnings, and occupational mobility

Table 3 reports results from estimating equation (2) on our main sample of workers born 1949-1960, using various sets of controls. Panel A shows that workers starting out in subsequently declining occupations in 1985 spent about nine months (0.73 years) less in employment than other workers over the following 28 years (column (1)). But this difference reduces to about six months when controlling for demographic characteristics and 1985 earnings (column (2)); this difference amounts to about two percent of the sample mean (23.4 years). In the most restrictive specification we compare those who experienced unanticipated occupational decline to observationally similar workers in similar occupations and industries, and this specification suggests that the losses from occupational decline were on average just over two months (0.2 years) of employment, which is less than one percent of the sample mean.

Panel B of Table 3 contains results from using cumulative earnings 1986-2013 as the outcome. Starting out in a later declining occupation was associated with 350,000 Swedish Krona (SEK) lower cumulative earnings, or about 5 percent of the sample mean.[37] Column (2) shows that controlling for individual characteristics, including 1985 earnings, the loss is very similar, though the confidence interval is much tighter. This is reasonable when we recall that declining and non-declining occupations were similar in terms of workers' 1985 earnings, as differences in gender and educational attainment offset each other. Further controlling for occupation-level life cycle profiles and growth predictors suggests an earnings loss of less than 2 percent of the sample mean, and this figure is again quite similar when we include occupation and industry dummies (columns (4)-(6)). Losses exceeding 3.5 percent of mean cumulative earnings lie outside the 95-percent confidence interval.

In panel C we examine earnings losses from occupational decline using an alternative earnings measure: cumulative earnings divided by predicted initial earnings (see Section 3.1.4 for details on the construction of this variable). With controls, 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.[38]

In Figure 2, we present a dynamic counterpart to the results reported in panel B, columns (2) and (6) of Table 3. Instead of cumulative earnings 1986-2013 as the outcome, we use cumulative earnings from 1986 up to the year indicated on the horizontal axis of each chart. The top panel of Figure 2 is suggestive of a smooth process of occupational decline, with earnings losses building up gradually. The picture is similar when we divide the coefficients by the mean of cumulative earnings at each horizon (bottom panel).

Next, we investigate how the probability that individuals remained in their starting occupations differed for those in declining occupations. Table 4 reports estimates of equation (2) with measures of

[37]We 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.

[38]Below we discuss results using alternative functional forms for cumulative earnings.

working in 2013 in the same occupation as in 1985 (or in a similar one) as the outcome. It is important to emphasize that 'not remaining' in the same occupation could reflect either occupational switching or non-employment, a point to which we return below.

Column (1) of panel A of Table 4 shows that without controls, 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 those in declining occupations to those with similar individual characteristics, occupational decline appears to reduce the probability of remaining in the 1985 occupation by 11 percentage points, and when we compare them to those with similar individual, occupation, and industry characteristics the estimate falls to 4.5 percentage points. Panels (B) and (C) of Table 4 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.[39]

So far we have focused on our declining variable, based on a 25-percent cutoff, which conservatively captures the occupations whose (US) employment fell over the period. We now explore specifications using alternative cutoffs, and report the results in Table A1. When we consider only declines of more than 50 percent, the employment and earnings losses are quite similar to those in the baseline case. However, in this specification the probability that workers remain in their starting occupation is reduced much more by occupational decline, with a range of 10-18 percentage points instead of the 4.5-11 percentage points in the baseline case. When we broaden the definition of declining to all occupational declines (using zero as the cutoff) the losses of employment and earnings are considerably smaller and in most cases imprecise, while the probability of remaining in the starting occupation is still significantly reduced by occupational decline. And when we broaden the definition to include all occupations whose growth was below the median (across Swedish workers in 1985) almost all the estimates are small and indistinguishable from zero. In sum, we tend to see larger consequences of occupational change for employment, earnings, and occupational switching when we focus on more restrictive definitions of occupational decline. Taken together, these results provide support for our definition of occupational decline as a way to isolate substantial adverse demand shocks occurring at the level of occupations, while also ensuring that

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

a meaningful fraction of workers is exposed to these shocks.[40] Finally, in the bottom panel of Table A1 we check the sensitivity of our results to excluding occupations with above-median growth form the control group. The last panel in Table A1 shows that our results remain virtually unchanged. This reassures us that our results are not driven by a comparison between declining and fast-growing occupations.

We provide further evidence on the relationship between our key outcomes of interest and US occupational employment growth in Figure A1. This figure shows outcome variables residualized based on the regression models in columns (2) and (6) in Tables 3 and 4, but with the declining dummy times its coefficient added. The mean difference in these transformed outcomes between declining and non-declining occupations is thus exactly equal to the coefficients reported in the tables. The outcomes thus transformed are then aggregated (averaged) to the level of 5-digit Swedish occupations that workers chose in 1985, and plotted against the percentage employment change 1984-2016 in the corresponding US occupations. As the figures show, the relationship between US occupational change and cumulative employment and earnings is almost flat, except for a drop in outcomes for those with substantial occupational decline. The probability of remaining in the starting occupation shows at best a weak positive relationship with US growth, although the probability of remaining is clearly smaller on average for the occupations classified as declining. Figure A2 shows a similar exercise except that we aggregate the data to the level of 3-digit (rather than 5-digit) Swedish occupations and put actual Swedish employment growth 1985-2013 (log changes) on the horizontal axis. This is the same variable that appeared on the left-hand side in the regressions reported in Table 1. Recall that the Swedish 3-digit occupations are balanced between 1985-2013 and thus cells are never empty, allowing us to plot log changes instead of percentage changes. When exploring bivariate relationships, log changes are preferable as their distribution is less skewed. In any case, the patterns in Figure A2 are qualitatively similar to those in Figure A1.

In Table A2, we report results from regressions using the full variation in occupational employment growth displayed in Figures A1 and A2. This variation likely confounds shocks to both supply and demand, which complicates the interpretation of these results, but we nevertheless report them for the sake of completeness. When we use the percentage employment change in US occupations corresponding to a worker's 1985 Swedish 5-digit occupation as an explanatory variable, there appear to be no differences between slow- and fast-growing occupations in terms of cumulative employment or the likelihood of re-

[40]Recall that in 1985, 13 percent of employees worked in subsequently declining occupations, under our default 25-percent cutoff. Using the 50-percent cutoff, the fraction of workers exposed falls substantially to only 6 percent. This also suggests that the changes in composition of workers non-declining occupations across most specifications above are fairly moderate.

maining in the initial occupation, but fast-growing occupations seem to have been associated with higher cumulative earnings. As can be expected given the skewed distribution of percentage changes, some of these results are quite sensitive to winsorizing this variable at the 95th percentile. When we use the log change in Swedish occupational employment, in the 3-digit occupation that workers chose in 1985, as an explanatory variable, we see positive associations of employment growth with cumulative earnings and with the probability of staying in the initial occupation, but no discernable correlation with cumulative employment.

Having found that employment and earnings losses due to occupational decline appear quite mild, we may worry that our empirical approach masks substantially larger losses. This may be because workers leaving declining occupations flock to non-declining occupations that are similar and that this increased labor supply depresses 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 address these concerns, Tables A3 and A4 report results from what we refer to as 'doughnut' specifications, namely the same regressions as those we report in Tables 3 and 4 but this time excluding 3-digit (SSYK96) occupations in which some but not all 5-digit occupations are declining. In Table A3 the coefficients are somewhat larger than in the baseline tables, but still the earnings losses vary from under 6 percent of mean earnings (rather than 5 percent) with individual controls to around 3 percent (rather than 2 percent) in the specification with all the controls. The coefficients on the probability of remaining in the starting occupation are also larger in magnitude in the specifications that we estimate in Table A4. In sum, these specifications suggest slightly larger losses from occupational decline, but they do not overturn the conclusion that losses appear to be mild, at least on average.

4.2 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 report results from estimating different versions of equation (3) in Table 5. 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, they do not vary when adding occupation and industry controls (including predictors of occupational change; even-numbered columns) over individual-level controls (odd-numbered columns), though the main coefficients on the declining dummy—giving the employment and earnings loss for the median worker—are sensitive to controls, consistent with our findings in the previous sub-section. 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).

These conclusions are robust to an alternative specification, whereby we 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 5 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 moderate mean losses reported in the previous sub-section mask more substantial losses for low earners (within an occupation).[41]

The pattern for the probability of remaining in the initial occupation appears to be non-monotonic. Linear interaction terms are of small magnitude and imprecisely estimated (panel A, columns (7)-(8)). However, when we switch to the more flexible specification, it appears that among the workers starting out in later declining occupations, both bottom-tercile and top-tercile workers were less likely to stay (panel B, columns (7)-(8)). These interaction coefficients are large at greater than ten percent of the overall mean (but 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 will discuss below.

One potential challenge in interpreting the results of Table 5 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 A5 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.

[41] This result bears similarities to Autor, Dorn, Hanson, and Song (2014), who find that low wage workers suffer more from trade-induced employment disruptions.

We briefly address a further dimension of heterogeneity; namely, we 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 (2) 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 A6 shows that among all workers, those who remained in their initial occupation had higher cumulative earnings, though in panel B we restrict the sample to those who were employed in 2013, and the finding reverses.[42] 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 4.7 below.

Before we investigate further outcomes, we check whether our findings concerning earnings losses—moderate losses on average, but more substantial losses for low-ranked workers—are robust to considering alternative functional forms of earnings. In panel A of Table A7 we consider discounted cumulative earnings, including a normalization by initial predicted earnings, applying a 5-percent discount rate. The results are qualitatively unchanged. Of course, given that losses from occupational decline appear to build gradually over time (Figure 2), and that a discounted sum puts more weights on early years, we estimate smaller losses from occupational decline when using earnings measures involving discounting. However, expressed in percent of the mean, the coefficients in panel A are only slightly lower than our estimates involving non-discounted cumulative earnings.[43] We further consider the percentile rank in cumulative earnings, the log of cumulative earnings, and the percentage change in earnings 1985-2013 as outcome measures. Again, the results remain qualitatively unchanged.

4.3 Unemployment, retraining, and early retirement

A natural question at this stage is to what extent we can attribute the losses in years of employment to unemployment and retraining spells, which, as discussed above, are available for the final 22 of the 28 years of our study. Table 6 reports estimates using the main specifications from Tables 3 and 5 but this time using cumulative days of unemployment (panel A) and state-sponsored retraining (panel B)

[42] Workers 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.

[43] In terms of discounted cumulative earnings, the average loss ranges from 1.5-4.5 percent of the overall mean depending on controls, and the bottom-tercile loss is about 5.5 percent of the bottom-tercile mean when including all controls.

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 most of these estimates 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 are however quantitatively small relative to the estimated employment losses. For bottom-tercile workers, we conservatively estimate an employment loss of 1.16 years.[44] Of these, unemployment and retraining account for only 22 percent.[45] 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 7 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 oc-

[44]From panel B, column (2) in Table 5 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.

[45]Of the mean employment loss, unemployment and retraining explain just over 40 percent, both when using individual controls and when including the complete set of controls. Recall however that the mean employment loss is statistically indistinguishable from zero when all controls are used (panel A, column (6) in Table 3).

cupational earnings rank, suggesting once more that those who earned least within their occupation to begin with suffered larger losses from occupational decline.

Panel B of Table 7 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.

4.4 What might workers' earnings have been in absence of occupational decline?

Our identification strategy assumes that conditional on detailed initial characteristics, outcomes for those in declining occupations would have looked similar to those in the non-declining occupations. While it is impossible to observe a counterfactual world without occupational decline, we present here some evidence that is consistent with the notion that workers in non-declining occupations constitute a reasonable counterfactual (after conditioning on observable characteristics).

First, we revisit Figure 2. The results of this figure suggest that those who were in declining occupations in 1985 did not suffer statistically significant earnings losses from 1985-1990. Admittedly, the sub-figures with only individual controls suggest that there may have been marginally significant losses over this period, but those figures with the full set of controls show a tightly estimated loss of close to zero over this period.

Second, the earnings losses for older workers reported in panel B of Table 7 are small, and when we include all the controls there is no indication of any loss. This is likely because many of these workers retire before the full impact of the occupational decline materializes.

Finally, we use data from 1975 and 1980 to study how earnings in the period before 1985 are correlated with our measure of declining occupations, conditional on our usual set of controls. The outcomes are earnings in 1980 (for those of our main group who were at least aged 25 in 1980) and the mean of 1975 and 1980 for middle-aged and older groups of cohorts. We display the results in Figure A3, where all regression coefficients are scaled by the in-sample mean of the outcome variable to ensure comparability. Those workers whose occupations would eventually decline did not appear to have significantly different earnings than others before the onset of the decline.[46]

Taken together, this evidence suggests that both before 1985 and in the years shortly after 1985, the workers in declining occupations did not significantly differ from those in non-declining occupations in

[46]We obtain these earnings data from the 1975 and 1980 population census. Earnings data for the population of Swedish workers are not available at annual frequency prior to 1985.

terms of earnings, after conditioning on observed characteristics. This supports the notion that workers in non-declining occupations form a reasonable control group (again, after conditioning on observed characteristics) and makes it more likely that the earnings losses we document were indeed a consequence of occupational decline.

4.5 Technology-related occupational decline

Our findings on occupational decline paint a consistent picture revealing moderate losses on average, but larger losses for workers with low initial earnings relative to others in their occupation. Consistent with much of the literature (Goos, Manning, and Salomons, 2014) we see technological change as a key driver of occupational decline, and especially occupational decline that is common to the US and Sweden. Nevertheless, there could be other drivers of occupational change, including changes one the supply side (changes in the workforce, trade shocks, or changes in government policy) and in final demand (such as taste changes). 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 3.1.2.[47]

In Table A8, we return to the level of 3-digit (SSYK96) occupations and relate aggregate employment changes to our measure of occupational decline as in Table 1, but this time also considering technology-linked declines. Employment changes in technology-declining occupations were no different, in a statistical sense, from the changes in other declining occupations (columns (3)-(4)). Compared to non-declining occupations, technology-declining occupations grew by 49 log points less, on average, once we control for other predictors of employment growth in column (6). This figure is similar to the results for the full set of declining occupations reported in Table 1.

Turning to the individual-level analysis,[48] Table A10 shows that 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—the coefficients on the technology-declining dummy in columns (1)-(2) are statistically indistinguishable from zero. Considering technology-related declines on their own (columns (3)-(5)), 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.

[47] Some 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.

[48] Workers 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 A9.

4.6 Studying occupational decline in the US using NLSY data

Our finding that the mean earnings and employment costs of occupational decline were limited (at least on average) raises the question whether this result is unique to Sweden and its institutional setting, or more general. We begin to address this question using 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 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 people 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 their 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 appendix.

To shed light on how occupational decline shaped earnings in the US, Table A11 reports estimates of specifications similar to those in panel B of Table 3. In panel A, we focus on raw mean earnings, which we compute for workers who report earnings for at least eight years (following Dahl and Lochner, 2012). Imposing only this requirement means the sample size remains large. Specifically, by not requiring observations in the later years, when response rates are lower, we make more use of the available data. The point estimates suggest that on average those in declining occupations earned almost exactly the same as the others once we control for individual characteristics and when we add all the further controls.[49] While the estimates are quite imprecise, the 95-percent confidence intervals exclude losses of 7 percent or more. We note 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. In panel B we repeat the exercise from panel A but this time using only odd years (since we do not have data on income in even years after 1994) and the picture is largely unchanged. In panel C we use data from all years but with imputations wherever possible, again following Dahl and Lochner (2012), as we explain in the appendix. Once again the results suggest no loss from occupational decline, although the precision is lower than before. In panel D we use cumulative earnings (instead of mean earnings) for those whose earnings are either reported or

[49] However, since we do not have data on income in even years after 1994, the earnings included are skewed to the beginning of the period we study, where the full effects of occupational decline may not yet be realized.

imputed for all years, and the estimate is again close to zero and imprecise. Finally, in panel E we divide the outcome variable from panel D by initial predicted income and here the point estimates are more negative, but again imprecise.

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).

Figure A4 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.

In Table A12 we repeat the analysis of Table A11 but this time using weeks employed and unemployed as outcomes. The 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). While the estimates are mostly imprecise, the 95-percent confidence intervals exclude losses of 2 percent or more, which again suggests that the moderate employment losses from occupational decline that we find in Sweden may generalize to the US setting.

Finally, in Table A13 we repeat the analysis from the previous two tables using the probability that a worker remains in their starting occupation as an outcome. Panel A reports estimates using 3-digit occupations, and the results suggest that occupational decline made it less likely that workers remained in their starting occupation. Estimates in the bottom two columns use broader occupational classifications (occupational group and broad occupational group) and are unfortunately imprecisely estimated.

The bottom line from the US analysis is, however, that the consequences of occupational decline in the US seem on average to be similar to those in Sweden, in that both lead to moderate mean employment or earnings losses, even though the US (NLSY) estimates are less precise.

4.7 Interpreting our findings through the lens of the theoretical model

Here we 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 modest earnings losses in declining occupations were associated with significant outflows from these occupations suggests that this mechanism may be relevant in our context. At the same time, our finding that workers with lower initial occupational earnings rank suffered larger losses suggest that even in Sweden there are real losers from occupational decline.

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 a occupational switching costs that decrease in workers' destination occupations, 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 involuntary 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 mobility responses to an adverse demand shock may differ depending on whether the shock was anticipated. We find that unanticipated declines are generally associated with a smaller mobility response, smaller aggregate employment declines, and also smaller earnings losses. While different combinations of model elements may give rise to different predictions, a possible interpretation of our findings is that the adverse demand changes behind unanticipated declines are of a smaller magnitude, as larger changes may be easier to foresee.

5 Conclusion

In this paper, we study the long run employment and earnings losses for individual workers when demand for their occupation 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 45 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 in declining occupations experienced relatively moderate earnings (employment) losses of 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, also suggests that mean employment and earnings losses were moderate.

While mean earnings losses from occupational decline were moderate, the losses for those at the bottom tercile of their occupations' earnings distribution in Sweden were larger (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.

Taken together, our findings suggest that the mean losses of occupational decline might be lower than feared. 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, governments should still be alert to the potential risks of occupational decline when technological capacity develops rapidly, 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 was mostly concentrated in low- to middle-skill occupations, whereas new technologies may displace high-skilled professionals. 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 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.

ACEMOGLU, D., AND P. RESTREPO (2018): "Modeling Automation," *AEA Papers and Proceedings*, 108, 48–53.

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 (2016): "The Risk of Automation for Jobs in OECD Countries: A Comparative Analysis," OECD Social, Employment and Migration Working Papers 189, OECD Publishing.

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., AND A. SALOMONS (2018): "Is Automation Labor ShareDisplacing? Productivity Growth, Employment, and the Labor Share," *Brookings Papers on Economic Activity*.

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., 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.

AUTOR, D. H., F. LEVY, AND R. J. MURNANE (2003): "The Skill Content Of Recent Technological Change: An Empirical Exploration," *The Quarterly Journal of Economics*, 118(4), 1279–1333.

BESSEN, J. (2016): "How computer automation affects occupations: Technology, jobs, and skills," Law and economics research paper 15-49, Boston University School of Law.

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.

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.

——— (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.

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.

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.

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.

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-," .

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.

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.

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.

LINDBECK, A. (1997): "The Swedish Experiment," *Journal of Economic Literature*, 35(3), 1273–1319.

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.

RUBIN, D. (1974): "Estimating causal effects of treatments in randomized and nonrandomized studies," *Journal of Educational Psychology*, 66(5), 688–701.

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.

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.

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.

VIKSTRÖM, J., AND G. VAN DEN BERG (2017): "Långsiktiga effekter av arbetsmarknadsutbildning," Discussion Paper Rapport 2017:17.

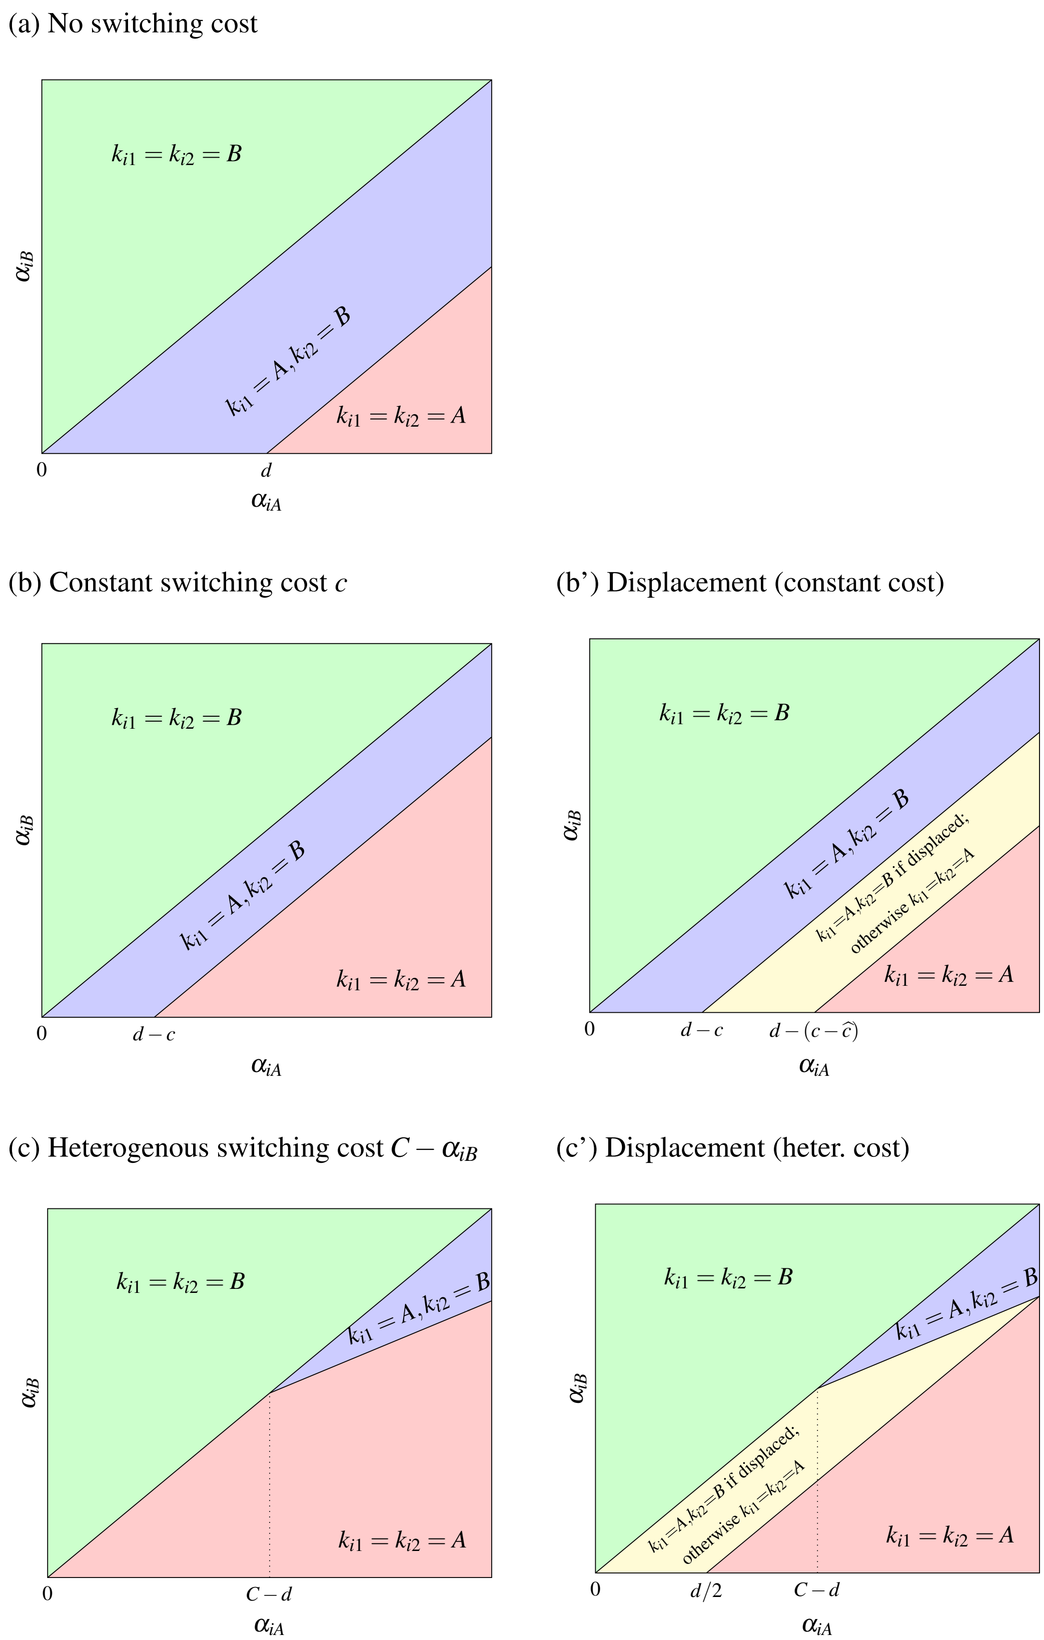

Notes: k_{it} denotes the occupation chosen by worker i in period t. α_{ik} denotes log productivity of worker i in occupation k. d is the amount by which the relative log occupational price declines from period 1 to period 2. The parameter values chosen are $(\overline{\alpha}, d, c, \widehat{c}, C)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

Figure 1: Sorting in a two-period Roy model

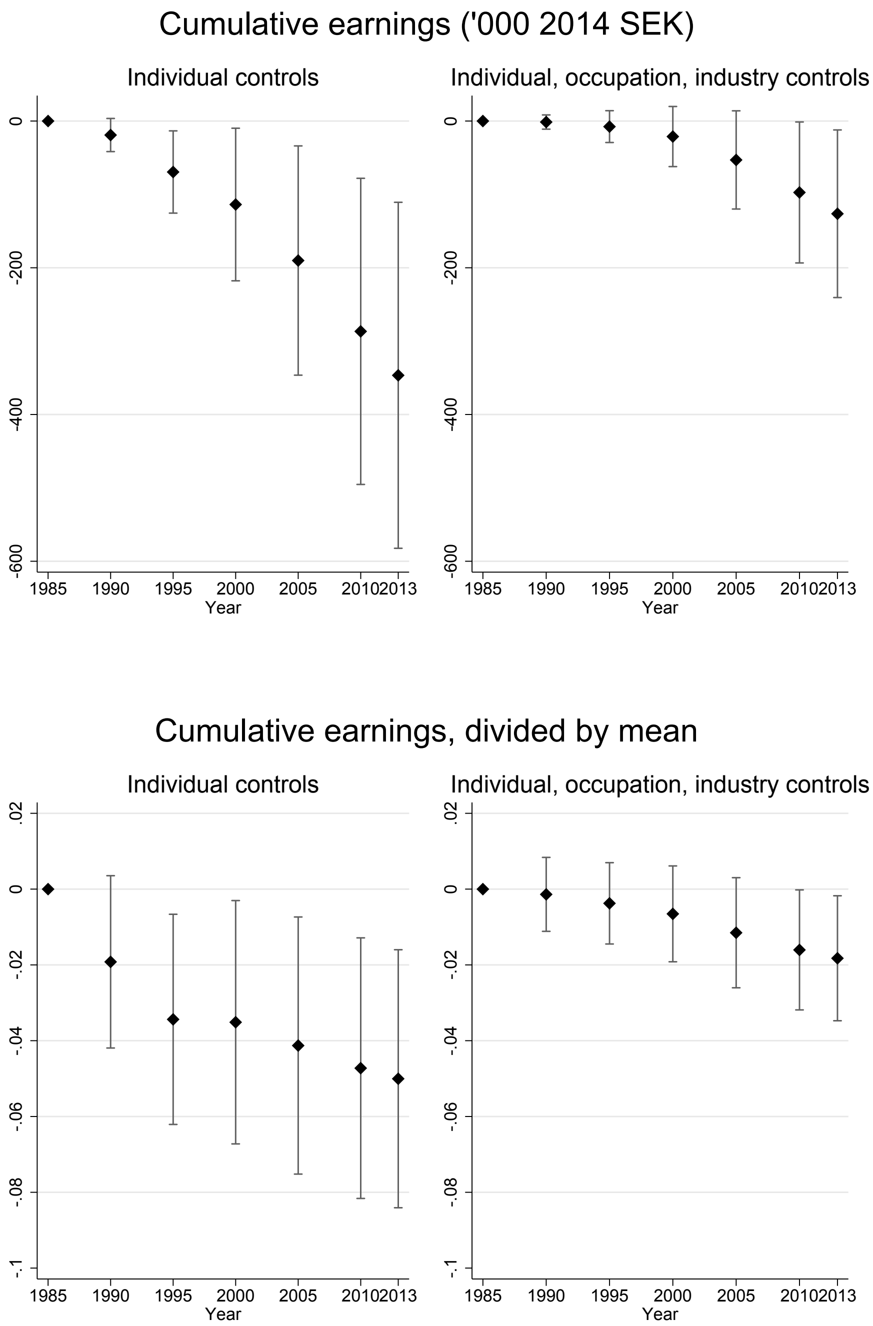

Notes: Diamonds mark the coefficients on the declining indicator from the regression specifications reported in panel B, columns (2) and (6) of Table 3, except that cumulative earnings are measured at various different times. Capped bars indicate 95-percent confidence intervals.

Figure 2: Occupational decline and cumulative earnings over time

Table 1: Employment growth in Swedish 3-digit occupations 1985-2013

	(1)	(2)	(3)	(4)	(5)	(6)
Declining	-0.76				-0.44	-0.46
	(0.17)				(0.18)	(0.18)
Employment share 1985		-1.23			-2.40	-2.31
		(1.61)			(1.57)	(1.53)
Employment growth 1960-85		0.34			0.16	0.15
		(0.08)			(0.09)	(0.08)
Predicted growth index			0.31		0.22	
			(0.07)		(0.08)	
Prediction: no change				-0.05		0.09
				(0.44)		(0.42)
Prediction: increase, slow				0.46		0.25
				(0.36)		(0.31)
Prediction: increase, average				0.74		0.55
				(0.29)		(0.25)
Prediction: increase, fast				1.13		0.82
				(0.29)		(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 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
A. Workers aged 16-64							
Intercept	0.52	39.5	0.33	0.56	0.11	191.3	0.25
	(0.078)	(0.41)	(0.030)	(0.033)	(0.027)	(10.8)	(0.050)
Declining	-0.25	-0.89	0.13	-0.063	-0.070	-0.23	0.38
	(0.088)	(0.63)	(0.035)	(0.034)	(0.028)	(11.0)	(0.085)
B. Workers aged 25-36							
Intercept	0.51	30.8	0.23	0.64	0.13	182.8	0.23
	(0.078)	(0.078)	(0.022)	(0.033)	(0.032)	(9.28)	(0.050)
Declining	-0.26	-0.19	0.15	-0.065	-0.082	12.0	0.38
	(0.085)	(0.091)	(0.030)	(0.034)	(0.034)	(9.40)	(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: Occupational decline and individual-level cumulative employment and earnings 1986-2013

	(1)	(2)	(3)	(4)	(5)	(6)
A. Cumulative years employed 1986-2013 (mean: 23.4)						
Declining	-0.73	-0.49	-0.49	-0.30	-0.24	-0.19
	(0.26)	(0.20)	(0.20)	(0.20)	(0.18)	(0.14)
B. Cumulative real earnings ('000 2014 SEK) 1986-2013 (mean: 6,926)						
Declining	-354	-347	-241	-117	-63	-126
	(419)	(120)	(81)	(76)	(71)	(58)
C. Cumulative real earnings divided by predicted initial earnings (mean: 38.7)						
Declining	-4.29	-2.10	-2.21	-1.52	-0.98	-1.11
	(0.91)	(0.53)	(0.54)	(0.54)	(0.41)	(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 3. Robust standard errors, clustered by 1985 3-digit occupation, in parentheses.

Table 4: Occupational decline and individual occupational stability

	(1)	(2)	(3)	(4)	(5)	(6)
A. Probability of working in same 3-digit occupation in 2013 as in 1985 (mean: 0.29)						
Declining	-0.14	-0.11	-0.11	-0.065	-0.086	-0.045
	(0.043)	(0.041)	(0.042)	(0.032)	(0.035)	(0.020)
B. Probability of working in same 2-digit occupation in 2013 as in 1985 (mean: 0.35)						
Declining	-0.12	-0.088	-0.087	-0.051	-0.070	-0.037
	(0.034)	(0.034)	(0.035)	(0.030)	(0.030)	(0.019)
C. Probability of working in same 1-digit occupation in 2013 as in 1985 (mean: 0.40)						
Declining	-0.098	-0.070	-0.069	-0.039	-0.060	-0.031
	(0.030)	(0.031)	(0.032)	(0.029)	(0.027)	(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 3 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 3, 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 5: Heterogeneity by within-occupation earnings rank

	Employment		Earnings		Earnings, normalized		Remain	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
A. Linear interaction								
Declining	-0.51	-0.23	-353.5	-131.0	-2.16	-1.19	-0.11	-0.045
	(0.21)	(0.15)	(110.7)	(55.8)	(0.55)	(0.37)	(0.041)	(0.020)
Declining × rank	1.17	1.17	441.5	449.2	2.63	2.63	-0.011	-0.0010
	(0.34)	(0.30)	(142.3)	(146.8)	(0.58)	(0.57)	(0.023)	(0.017)
B. Dummy interactions								
Declining	-0.32	-0.031	-323.2	-98.0	-1.94	-0.97	-0.083	-0.022
	(0.24)	(0.18)	(123.8)	(66.7)	(0.54)	(0.41)	(0.045)	(0.021)
Declining × bottom tercile	-1.12	-1.13	-341.8	-350.1	-2.10	-2.06	-0.046	-0.040
	(0.35)	(0.33)	(106.7)	(101.5)	(0.54)	(0.51)	(0.014)	(0.013)
Declining × top tercile	0.54	0.55	232.3	235.1	1.37	1.40	-0.047	-0.030
	(0.20)	(0.16)	(135.8)	(132.1)	(0.43)	(0.48)	(0.027)	(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 3, and for columns (7)-(8) it is the same as that in Table 4. Robust standard errors, clustered by 1985 3-digit occupation, in parentheses.

Table 6: Occupational decline and the incidence of unemployment and retraining

	Ever				Cumulative days			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
A. Unemployment								
Declining	0.041	0.013	0.015	0.019	52.4	17.9	20.8	20.5
	(0.021)	(0.013)	(0.012)	(0.015)	(24.8)	(14.0)	(14.0)	(18.2)
Declining × rank			-0.036				-63.8	
			(0.012)				(21.5)	
Declining × bottom tercile				0.017				42.4
				(0.012)				(18.3)
Declining × top tercile				-0.033				-43.7
				(0.012)				(17.0)
Mean of dep. var.		0.39				262		
Mean of dep. var., bottom		0.43				317		
B. Retraining								
Declining	0.035	0.012	0.013	0.015	11.4	4.73	5.04	5.81
	(0.010)	(0.0064)	(0.0063)	(0.0081)	(2.68)	(1.46)	(1.48)	(2.26)
Declining × rank			-0.027				-8.63	
			(0.0070)				(1.98)	
Declining × bottom tercile				0.014				4.38
				(0.0072)				(2.28)
Declining × top tercile				-0.022				-6.96
				(0.0064)				(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 Table 3. See the notes to Table 5 for a description of right-hand side variables. Robust standard errors, clustered by 1985 3-digit occupation, in parentheses.

Table 7: 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	-0.32	-0.47	-273.1	-72.9	-99.4	-0.39	-0.15	-0.25
	(0.16)	(0.11)	(0.12)	(53.0)	(43.4)	(39.7)	(0.097)	(0.065)	(0.074)
Declining × rank			0.98			173.6			0.65
			(0.25)			(85.8)			(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	-0.047	-0.087	-75.0	12.3	8.09	-0.19	-0.011	-0.038
	(0.085)	(0.070)	(0.072)	(18.2)	(18.8)	(18.2)	(0.062)	(0.048)	(0.049)
Declining × rank			0.18			14.4			0.13
			(0.093)			(26.4)			(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 5 for a description of right-hand side variables. Robust standard errors, clustered by 1985 3-digit occupation, in parentheses.

Appendix for Online Publication

Theory appendix

Here we give formal derivations proving the claims made in Section 2 of the paper.

Start with the baseline scenario of costless switching and an unanticipated price change from Section 2.2, and corresponding to Figure 1, panel (a). Recall that switchers must satisfy the inequalities $\alpha_{iB} \leq \alpha_{iA}$ and $\alpha_{iB} > \alpha_{iA} - d$. Panel (a) of Figure 1 indicates that there are two cases. Among workers in occupation A with $\alpha_{iA} < d$, 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 $\alpha_{iA} \geq d$, 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 $d/2$. Taken together, the probability of switching is $\mathbb{P}(k_{i2} = B | k_{i1} = A, \alpha_{iA}) = \mathbb{1}\{\alpha_{iA} < d\} + \mathbb{1}\{\alpha_{iA} \geq d\}(d/\alpha_{iA})$, which is weakly decreasing in α_{iA}; and the expected period-2 earnings losses of workers starting out in occupation A, conditional on their productivity in A, are given by

$$\mathbb{E}[\text{loss}|\alpha_{iA}] = \mathbb{1}\{\alpha_{iA} < d\}\frac{\alpha_{iA}}{2} + \mathbb{1}\{\alpha_{iA} \geq d\}d\left(1 - \frac{d}{2\alpha_{iA}}\right).$$

The mean earnings loss is thus strictly increasing in initial earnings.

Now turn to the case of a constant switching cost from Section 2.3, and corresponding to Figure 1, panel (b). Choices are no longer period-by-period: period-1 choices affect expected utility in period 2. 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 $\mathbb{E}[\max\{\pi_{A2} + \alpha_{iA}, \pi_{B2} + \alpha_{iB} - c\}] = \mathbb{E}[\max\{\pi_{A2} + \alpha_{iA}, \pi_{B2} + \alpha_{iB}\}] = \alpha_{iA}$, and $V_{iA} \geq V_{iB}$.
- If $\alpha_{iB} - c \leq \alpha_{iA} < \alpha_{iB}$, then $\mathbb{E}[\max\{\pi_{A2} + \alpha_{iA}, \pi_{B2} + \alpha_{iB} - c\}] = \alpha_{iA}$, $\mathbb{E}[\max\{\pi_{A2} + \alpha_{iA}, \pi_{B2} + \alpha_{iB}\}] = \alpha_{iB}$, and $V_{iA} < V_{iB}$.
- If $\alpha_{iA} < \alpha_{iB} - c$, then $\mathbb{E}[\max\{\pi_{A2} + \alpha_{iA}, \pi_{B2} + \alpha_{iB} - c\}] = \alpha_{iB} - c$, $\mathbb{E}[\max\{\pi_{A2} + \alpha_{iA}, \pi_{B2} + \alpha_{iB}\}] = \alpha_{iB}$, and $V_{iA} < V_{iB}$.

This establishes that worker i chooses occupation A in period 1 if and only if $\alpha_{iA} \geq \alpha_{iB}$. After the price change in period 2, worker i switches if and only if $\alpha_{iB} - c > \alpha_{iA} - d$.

The same analysis can be applied to the case of a heterogenous switching cost $C - \alpha_{iB}$, with $C > \overline{\alpha}$, corresponding to Figure 1, panel (c). 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 $\mathbb{E}[\max\{\pi_{A2} + \alpha_{iA}, \pi_{B2} + \alpha_{iB} - (C - \alpha_{iB})\}] = \mathbb{E}[\max\{\pi_{A2} + \alpha_{iA}, \pi_{B2} + \alpha_{iB}\}] = \alpha_{iA}$, and $V_{iA} \geq V_{iB}$.
- If $\alpha_{iB} - (C - \alpha_{iB}) \leq \alpha_{iA} < \alpha_{iB}$, then $\mathbb{E}[\max\{\pi_{A2} + \alpha_{iA}, \pi_{B2} + \alpha_{iB} - (C - \alpha_{iB})\}] = \alpha_{iA}$, $\mathbb{E}[\max\{\pi_{A2} + \alpha_{iA}, \pi_{B2} + \alpha_{iB}\}] = \alpha_{iB}$, and $V_{iA} < V_{iB}$.

- If $\alpha_{iA} < \alpha_{iB} - (C - \alpha_{iB})$, then $\mathbb{E}[\max\{\pi_{A2} + \alpha_{iA}, \pi_{B2} + \alpha_{iB} - (C - \alpha_{iB})\}] = \alpha_{iB} - (C - \alpha_{iB})$, $\mathbb{E}[\max\{\pi_{A2} + \alpha_{iA}, \pi_{B2} + \alpha_{iB}\}] = \alpha_{iB}$, 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 $\alpha_{iB} - (C - \alpha_{iB}) > \alpha_{iA} - d$. Switching probabilities are thus $\mathbb{P}(k_{i2} = B | k_{i1} = A, \alpha_{iA}) = \mathbb{1}\{\alpha_{iA} > C - d\}(1/2 - (C-d)/(2\alpha_{iA}))$, an expression which is strictly increasing in α_{iA} when $\alpha_{iA} > C - d$. Mean period-2 earnings of switchers equal $3\alpha_{iA}/2 - (C+d)/2$, so their mean earnings loss is $-\alpha_{iA}/2 + (C+d)/2$, which is strictly decreasing in α_{iA}.

Now let us analyse the case with probabilistic displacement and heterogeneous reemployment costs from Section 2.4, and corresponding to Figure 1, panel (c'). Let the probability that workers in occupation A are displaced, $\lambda \in (0,1)$, be constant and thus independent of α_{iA}. This means that the overall switching probability is given by

$$\mathbb{P}(\text{switch}|\alpha_{iA}) = \mathbb{1}\{\alpha_{iA} \leq d/2\}\lambda + \mathbb{1}\{\alpha_{iA} \in (d/2, C-d]\}\frac{\lambda d}{2\alpha_{iA}} + \mathbb{1}\{\alpha_{iA} > C-d\}\left(\frac{1-\lambda}{2} - \frac{(1-\lambda)C - d}{2\alpha_{iA}}\right).$$

The switching probability is thus first constant, then strictly decreasing, then strictly increasing in initial earnings (the last property relies on the displacement probability not being too large, so that $(1-\lambda)C > d$, an assumption we shall maintain). But how do expected losses change with initial earnings? For $\alpha_{iA} \leq d/2$, 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 $\lambda C + (1-\lambda)d$. 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). In this region the overall expected loss is thus

$$\mathbb{E}[\text{loss}|\alpha_{iA}] = \lambda\left[C - \alpha_{iA} + d - \frac{d^2}{4\alpha_{iA}}\right] + (1-\lambda)d,$$

which is strictly decreasing in α_{iA} in this range. Finally, for $\alpha_{iA} > C - d$ there are some workers moving voluntarily, even if not displaced. Their fraction (earnings loss) is increasing (decreasing) in initial earnings, and moreover their earnings losses are less than those of non-displaced stayers, so that again, the overall earnings loss is decreasing in initial earnings.

If period-2 occupational prices are revealed at the start of period 1, 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 cots 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).

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.

Mapping of OOH occupations to Swedish occupations

Recall from Section 3.1.2 that 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. Here we discuss how this weighted employment growth is calculated based on the mapping from OOH occupations to Swedish YRKE5 occupations. We also explain how we identify technology-related declines.

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.[50]

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

$$g_s \equiv \frac{N_{s,2016} - N_{s,1984}}{N_{s,1984}}, \tag{A1}$$

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

$$g_s \equiv \underset{1\times K}{\alpha_s} \times \underset{K\times 1}{g}, \tag{A2}$$

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.[51] 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,

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

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

[51] Note 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 (A1) and (A2) is easily shown:

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

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

The computation of the total changes in equation (A1), or the weights and changes in equation (A2) 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 Molders 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

$$\hat{\alpha}_{s,343} \times \hat{g}_{343} + \alpha_{s,344} \times \hat{g}_{344}$$

$$= 0.988 \times (-0.286) + 0.012 \times (-1.00) = -0.295$$

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,

$$\alpha_{s,343} \times \hat{g}_{343} + \alpha_{s,344} \times \hat{g}_{344}$$

$$= 0.414 \times (-0.286) + 0.586 \times (-1.00) = -0.704$$

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,

$$\hat{g}_s = \alpha_s \times \hat{g}. \tag{A3}$$

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.[52]

- 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.

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

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).[53] 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 $\hat{g}_k < -0.25$.

[52] We 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".

[53] There 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:

$$\hat{g}_s \equiv \underset{1\times K}{\alpha_s} \times \underset{K\times K}{\mathbb{1}\{\text{technology}\}} \times \underset{K\times 1}{\hat{g}} \quad + \quad \underset{1\times K}{\alpha_s} \times \big(I - \underset{K\times K}{\mathbb{1}\{\text{technology}\}}\big) \times \underset{K\times 1}{\hat{g}}\,, \tag{A4}$$

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 (A4) is less than -0.25, formally

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

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)

Description of NLSY data

This appendix describes the data we use to study the individual consequences of occupational decline in the United States. 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.

Data sources

National Longitudinal Survey Youth 1979 (NLSY79)

The main dataset we use to study occupational decline and its consequences in the US is the NLSY79 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)[54] with updated weights to include only those in the sample as of 1987 (see below), and updates for recent errata from the NLSY.[55] 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.

1980 Census

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).

Income inflation

To convert income to $2014, we use the chained Consumer Price Index for All Urban Consumers (CPI-U), published by the BLS and made available by the Federal Reserve Bank of Minneapolis (2018).

[54] We thank Michael Böhm for generously sharing this data and his expertise.

[55] We 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)

Sample selection and weighting

Sample selection

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.

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.

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).

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.

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 criteria[56] 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

[56]The 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

the 2000 census codes, for all employers; we consider only the occupation associated with each person's main job.

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)").

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. 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 codes[57] 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.

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 mapped 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.

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

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 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).

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.

Income

The income 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.

Income measures

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 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 people-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 above.

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 to maintain a sample of people that is as representative as possible of the US population of the relevant age groups. 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.

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.

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 NLS 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 year, 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.

Specifically, 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% [58]
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

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 earning. Extrapolating more than two years may result in implausibly

[58] To 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.

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.

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.

As in the income regressions, we restrict the sample to those with at least 8 years (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.

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.

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.

Appendix figures and tables

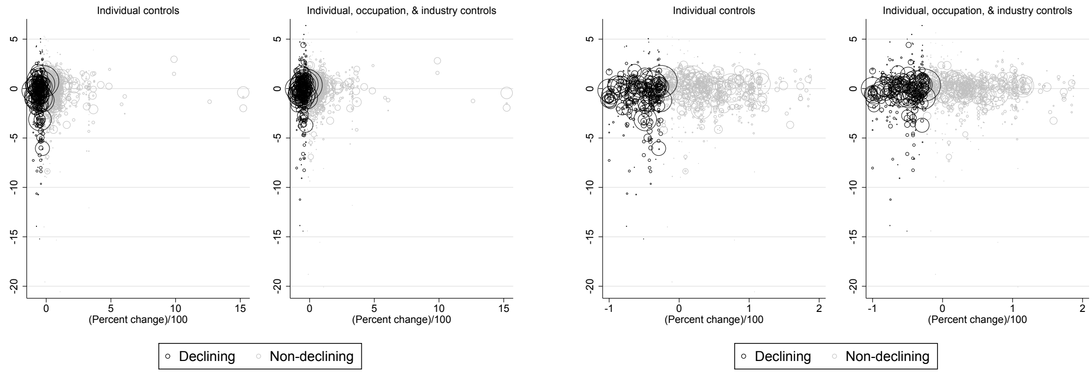

(a) Cumulative employment

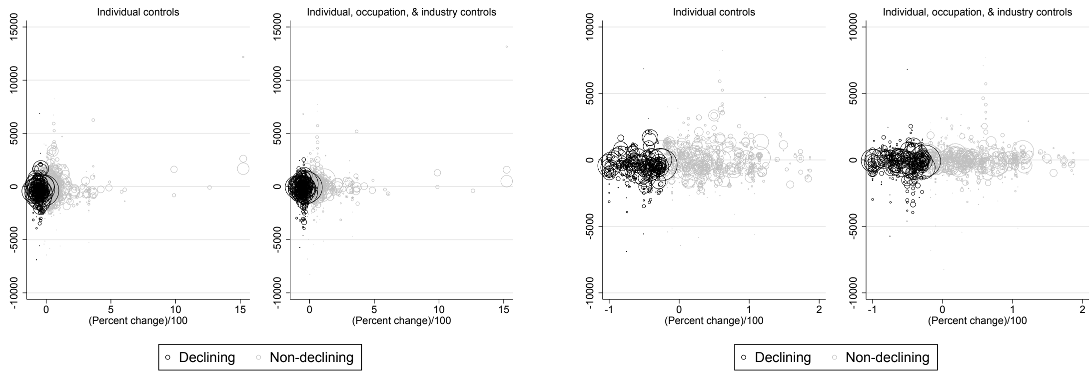

(b) Cumulative earnings

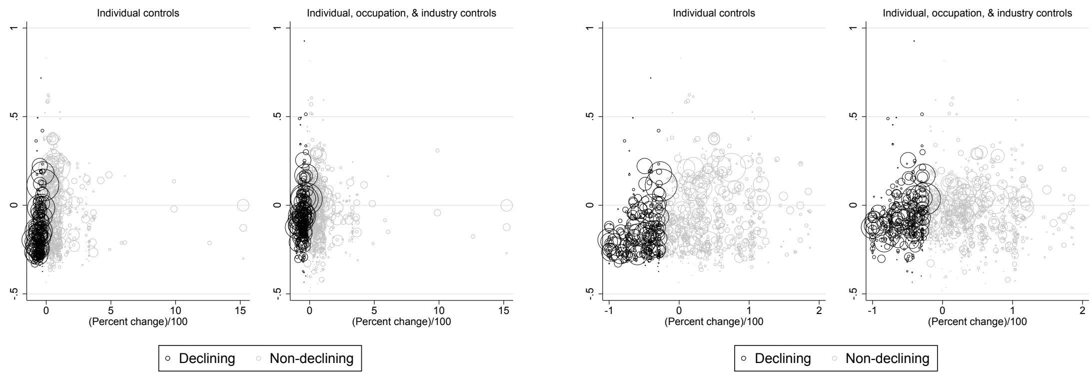

(c) Probability of remaining in the initial 3-digit occupation

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 occupations(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 3 and 4, 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 A1: Main outcomes and percent change in employment (US)

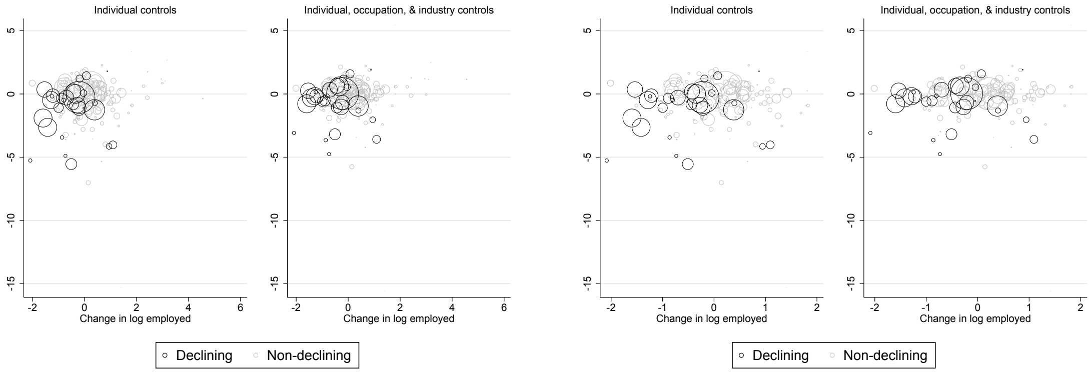

(a) Cumulative employment

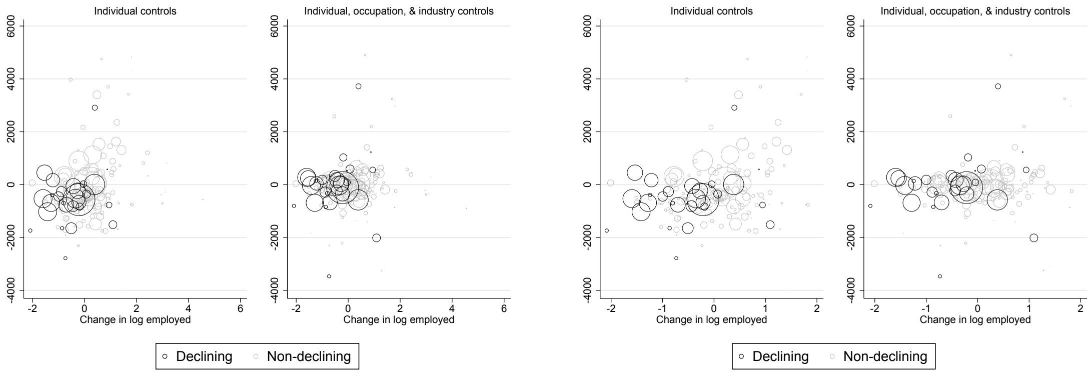

(b) Cumulative earnings

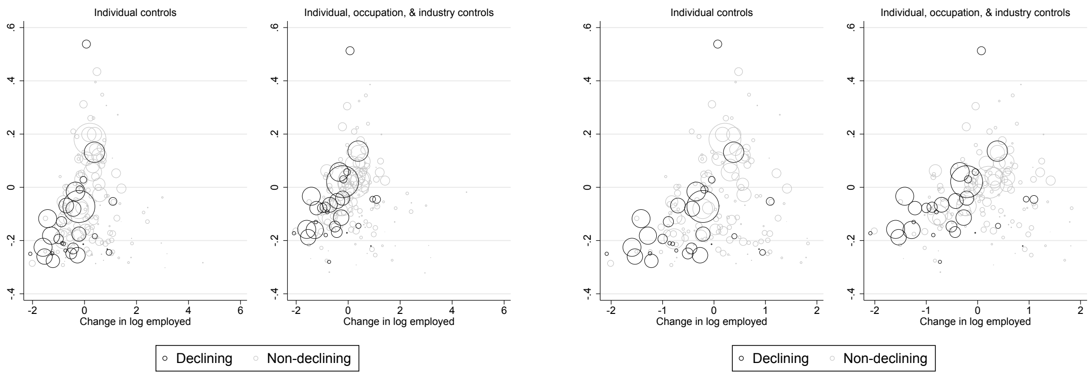

(c) Probability of remaining in the initial 3-digit occupation

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 A2, 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 A2). The pairs of graphs on the right are truncated versions of those on the left.

Figure A2: Main outcomes and change in log employment (Sweden)

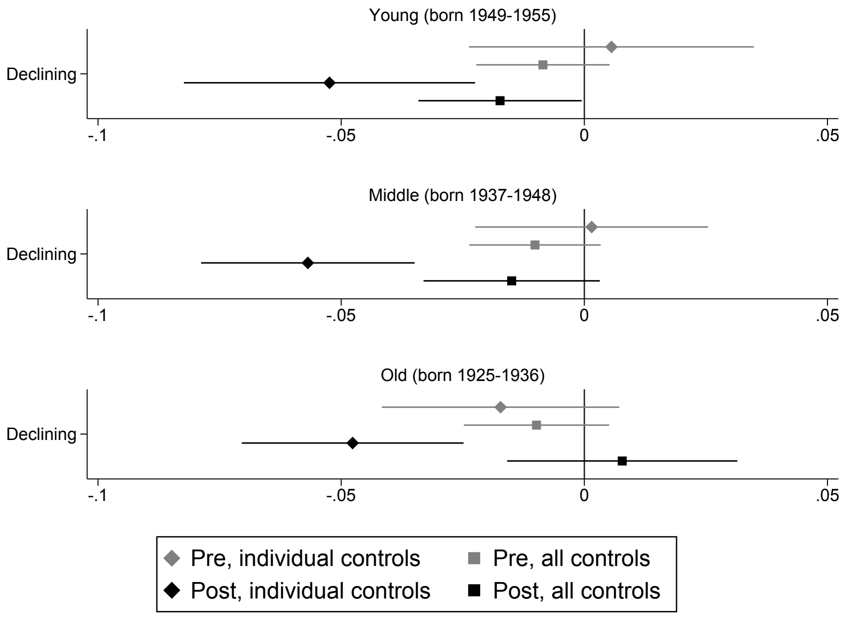

Notes: Coefficients on the declining indicator along with their 95-percent confidence intervals (robust to clustering by 1985 3-digit occupation) are displayed, where the regressions vary the sample, controls, and outcome variables. Coefficients are scaled by the mean of the outcome variable in each 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 3, and 'all controls' are the ones from column (6) in that table.

Figure A3: Earnings prior to occupational decline

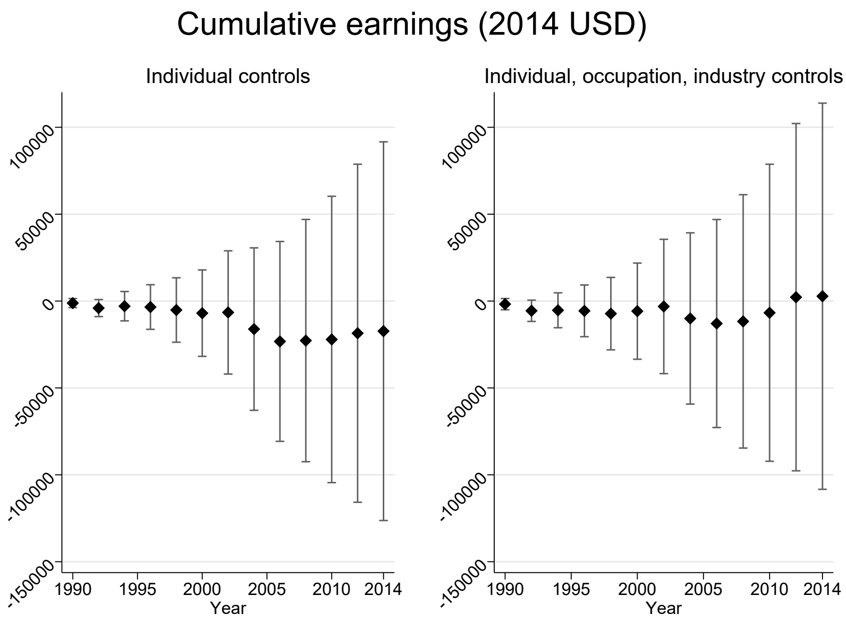

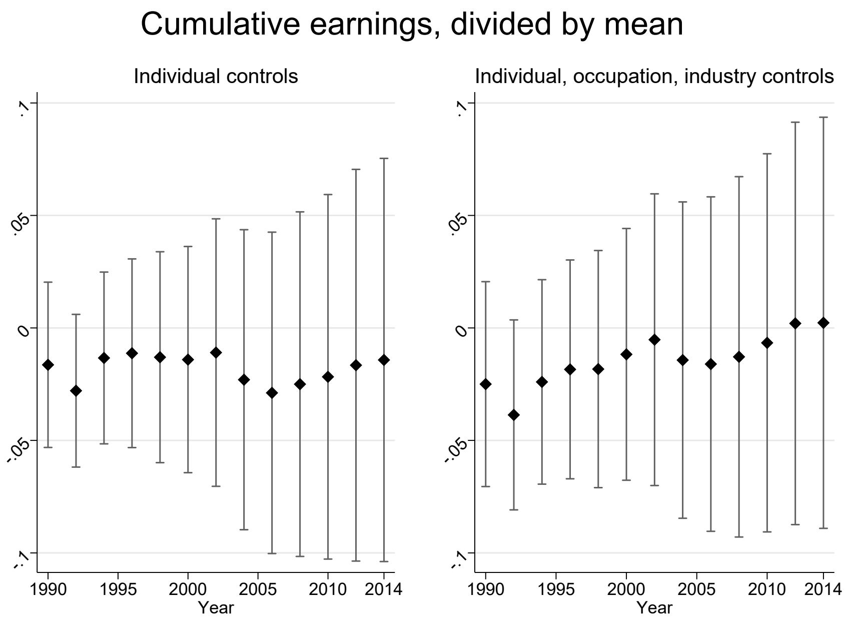

Notes: Diamonds mark the coefficients on the declining indicator from the regression specifications reported in columns (2) and (6) of Table A11, except that income accumulates only through time t. Capped bars indicate 95-percent confidence intervals.

Figure A4: US (NLSY) occupational decline and individual-level earnings over time

Table A1: 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	-0.18	-248.1	-90.0	-2.44	-0.98	-0.18	-0.10
	(0.20)	(0.15)	(115.6)	(75.7)	(0.62)	(0.43)	(0.040)	(0.020)
Percent change $\in [-100, -25)$ (baseline)	-0.49	-0.19	-346.6	-126.4	-2.10	-1.11	-0.11	-0.045
	(0.20)	(0.14)	(120.3)	(58.3)	(0.53)	(0.36)	(0.041)	(0.020)
Percent change $\in [-100, 0)$	-0.043	-0.0030	-35.0	-57.5	-0.70	-0.91	-0.15	-0.063
	(0.20)	(0.13)	(158.8)	(74.7)	(0.70)	(0.47)	(0.041)	(0.021)
Percent change $\in [-100, 31)$ (below median)	0.14	0.15	-46.5	-61.9	-0.55	-0.53	-0.087	-0.0094
	(0.18)	(0.13)	(150.7)	(76.1)	(0.57)	(0.50)	(0.037)	(0.022)
Baseline; control: percent change $\in (-25, 31)$	-0.72	-0.27	-460.5	-126.6	-2.40	-1.17	-0.077	-0.053
	(0.22)	(0.16)	(123.3)	(61.9)	(0.51)	(0.40)	(0.038)	(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 3 and 4 for further descriptions of variables and sample definitions. Robust standard errors, clustered by 1985 3-digit occupation, in parentheses.

Table A2: 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	-0.026	103.7	64.7	0.47	0.25	0.0058	-0.0020
	(0.037)	(0.036)	(30.2)	(14.9)	(0.11)	(0.13)	(0.0068)	(0.0029)
Percent employment change / 100 (US), winsorized	0.010	0.000027	83.8	91.1	0.86	0.46	0.051	0.0035
	(0.11)	(0.080)	(112.0)	(47.5)	(0.40)	(0.25)	(0.025)	(0.014)
Log employment change (SWE)	-0.034	0.049	306.4	73.7	0.85	0.087	0.11	0.066
	(0.15)	(0.11)	(135.1)	(65.9)	(0.50)	(0.50)	(0.031)	(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 3 and 4 for further descriptions of variables and sample definitions. Robust standard errors, clustered by 1985 3-digit occupation, in parentheses.

Table A3: Occupational decline and individual-level cumulative employment and earnings 1986-2013—'doughnut' specifications

	(1)	(2)	(3)	(4)	(5)	(6)
A. Cumulative years employed 1986-2013 (mean: 23.5)						
Declining	-1.46	-0.97	-0.97	-0.82	-0.35	-0.41
	(0.53)	(0.42)	(0.42)	(0.46)	(0.28)	(0.29)
B. Cumulative real earnings ('000 2014 SEK) 1986-2013 (mean: 6,612)						
Declining	-484	-403	-333	-140	-81	-217
	(608)	(196)	(177)	(181)	(158)	(167)
C. Cumulative real earnings divided by predicted initial earnings (mean: 39.2)						
Declining	-5.40	-2.49	-2.56	-1.81	-1.18	-1.69
	(1.33)	(1.09)	(0.98)	(1.07)	(0.82)	(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 3, 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 A4: Occupational decline and individual occupational stability—'doughnut' specifications

	(1)	(2)	(3)	(4)	(5)	(6)
A. Probability of working in same 3-digit occupation in 2013 as in 1985 (mean: 0.35)						
Declining	-0.25	-0.21	-0.21	-0.12	-0.17	-0.10
	(0.046)	(0.051)	(0.052)	(0.044)	(0.046)	(0.046)
B. Probability of working in same 2-digit occupation in 2013 as in 1985 (mean: 0.40)						
Declining	-0.21	-0.16	-0.16	-0.089	-0.12	-0.059
	(0.039)	(0.045)	(0.046)	(0.043)	(0.045)	(0.042)
C. Probability of working in same 1-digit occupation in 2013 as in 1985 (mean: 0.44)						
Declining	-0.19	-0.14	-0.14	-0.077	-0.11	-0.045
	(0.036)	(0.042)	(0.043)	(0.042)	(0.043)	(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 4, but with the 'doughnut' restrictions from Table A3 applied. The number of observations is 333,357. Robust standard errors, clustered by 1985 3-digit occupation, in parentheses.

Table A5: 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	-0.20	-332.3	-154.0	-2.32	-1.22	-0.11	-0.042
	(0.22)	(0.14)	(90.2)	(59.3)	(0.56)	(0.37)	(0.041)	(0.020)
Declining × rank	0.92	0.96	407.7	439.5	2.33	2.41	-0.020	-0.014
	(0.33)	(0.29)	(141.9)	(137.3)	(0.59)	(0.56)	(0.016)	(0.015)
B. Dummy interactions								
Declining	-0.26	0.048	-302.5	-94.5	-1.94	-0.92	-0.095	-0.032
	(0.22)	(0.16)	(96.4)	(62.9)	(0.52)	(0.38)	(0.050)	(0.025)
Declining × bottom tercile	-1.16	-1.11	-370.4	-390.7	-2.14	-2.10	-0.015	-0.0082
	(0.36)	(0.33)	(93.5)	(86.2)	(0.48)	(0.44)	(0.019)	(0.017)
Declining × top tercile	0.16	0.24	220.2	202.7	0.99	1.12	-0.037	-0.026
	(0.15)	(0.15)	(109.2)	(111.3)	(0.47)	(0.45)	(0.027)	(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 5 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 A6: Cumulative earnings of leavers and stayers in declining and non-declining occupations

	(1)	(2)	(3)	(4)
A. All workers (553,169 observations)				
Remain	335	303	305	284
	(122)	(91)	(133)	(101)
Declining			-272	-127
			(122)	(90)
Declining × remain			177	190
			(239)	(185)
B. Employed in 2013 (404,043 observations)				
Remain	-398	-498	-439	-531
	(115)	(66)	(124)	(72)
Declining			-357	-188
			(123)	(94)
Declining × remain			238	312
			(231)	(158)
C. Employed in 2013, bottom third (140,892 observations)				
Remain	-109	-285	-133	-307
	(139)	(85)	(143)	(80)
Declining			-418	-238
			(145)	(173)
Declining × remain			-32	235
			(596)	(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 4, except for the restrictions indicated. Sampling weights are applied. Robust standard errors, clustered by 1985 3-digit occupation, in parentheses.

Table A7: Alternative functional forms for earnings

A. Discounted cumulative earnings

	Discounted cumulative earnings				Discounted cumulative earnings, normalized			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Declining	-152.7	-47.8	-49.5	-33.2	-0.94	-0.47	-0.51	-0.40
	(57.1)	(25.5)	(24.4)	(29.5)	(0.25)	(0.16)	(0.16)	(0.18)
Declining × rank			213.9				1.22	
			(68.5)				(0.25)	
Declining × bottom tercile				-166.5				-0.96
				(47.6)				(0.23)
Declining × top tercile				109.3				0.64
				(62.0)				(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	-0.85	-0.85	-0.95	-0.060	-0.021	-0.026	-0.00054	-41.8	-11.7	-9.39	0.73
	(0.84)	(0.54)	(0.50)	(0.63)	(0.022)	(0.013)	(0.014)	(0.017)	(11.2)	(8.61)	(8.80)	(7.99)
Declining × rank			5.15				0.17				145.6	
			(0.93)				(0.035)				(35.7)	
Declining × bottom tercile				-3.26				-0.15				-110.0
				(0.78)				(0.037)				(32.7)
Declining × top tercile				3.41				0.072				69.9
				(0.89)				(0.018)				(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 3 and 4 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 A8: Employment growth in Swedish 3-digit occupations 1985-2013—technology-related declines

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

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 A9: 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
A. Occupational decline, pooled							
Intercept	0.52	39.5	0.33	0.56	0.11	191.3	0.25
	(0.078)	(0.41)	(0.030)	(0.033)	(0.027)	(10.8)	(0.050)
Declining	-0.25	-0.89	0.13	-0.063	-0.070	-0.23	0.38
	(0.088)	(0.63)	(0.035)	(0.034)	(0.028)	(11.0)	(0.085)
B. Occupational decline, by presence of technology link							
Intercept	0.52	39.5	0.33	0.56	0.11	191.3	0.25
	(0.078)	(0.41)	(0.030)	(0.033)	(0.027)	(10.8)	(0.050)
Declining	-0.32	0.033	0.13	-0.086	-0.041	5.31	0.26
	(0.10)	(0.87)	(0.056)	(0.051)	(0.035)	(15.0)	(0.10)
Declining (technology)	0.11	-1.49	0.010	0.037	-0.047	-8.90	0.20
	(0.097)	(1.01)	(0.059)	(0.050)	(0.025)	(14.6)	(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 A8 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 A10: Occupational decline and individual-level cumulative employment and earnings 1986-2013—technology-related decline

	(1)	(2)	(3)	(4)	(5)
A. Cumulative years employed 1986-2013 (mean: 23.4)					
Declining	-0.93	-0.45			
	(0.44)	(0.24)			
Declining (technology)	0.72	0.42	-0.21	0.01	-0.16
	(0.45)	(0.23)	(0.16)	(0.14)	(0.12)
Declining (tech) × rank					1.31
					(0.36)
B. Cumulative real earnings ('000 2014 SEK) 1986-2013 (mean: 6,926)					
Declining	-426	-181			
	(232)	(93)			
Declining (technology)	128	87	-303	-107	-122
	(262)	(102)	(131)	(65)	(61)
Declining (tech) × rank					491
					(155)
C. Probability of working in same 3-digit occupation in 2013 as in 1985 (mean: 0.29)					
Declining	-0.077	-0.029			
	(0.051)	(0.022)			
Declining (technology)	-0.058	-0.025	-0.135	-0.053	-0.056
	(0.044)	(0.029)	(0.043)	(0.026)	(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 A8 for a description of these indicators). Columns (1)-(2) are based on the same samples as the results in Tables 3 and 4. 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 A11: US (NLSY) occupational decline and individual-level earnings 1988-2013

	(1)	(2)	(3)	(4)	(5)	(6)
A. Average real earnings (2014 USD) 1988-2013, no interpolation (mean: 44,083)						
Declining	-2,661	-151	279	595	655	-24
	(3,372)	(1,589)	(1,635)	(1,750)	(1,584)	(1,536)
B. Average real earnings (2014 USD) 1989-2013, no interpolation, odd years only (mean: 46,057)						
Declining	-2,600	123	607	954	1,124	384
	(3,649)	(1,901)	(1,967)	(2,107)	(1,891)	(1,823)
C. Average real earnings (2014 USD) 1988-2013 (mean: 46,891)						
Declining	-2,970	-92	408	892	963	227
	(3,783)	(2,029)	(2,100)	(2,255)	(2,012)	(1,969)
D. Cumulative real earnings (2014 USD) 1988-2013 (mean: 1,216,117)						
Declining	-95,964	-17,313	-5,801	12,310	23,322	2,783
	(102,583)	(55,596)	(56,336)	(59,240)	(54,961)	(56,695)
E. Cumulative real earnings divided by predicted initial earnings (mean: 44.2)						
Declining	-4.04	-2.71	-3.69	-4.87	-2.10	-2.43
	(2.57)	(2.44)	(2.10)	(3.45)	(2.80)	(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 A12: US (NLSY) occupational decline and individual employment

	(1)	(2)	(3)	(4)	(5)	(6)
A. Fraction of reported weeks in employment status (mean: 0.83)						
Declining	0.014	0.017	0.018	0.014	0.0035	0.0042
	(0.012)	(0.0091)	(0.0090)	(0.0091)	(0.0089)	(0.0094)
B. Fraction of reported weeks in unemployment status (mean: 0.03)						
Declining	0.0063	0.0043	0.0041	0.0052	0.0051	0.0042
	(0.0038)	(0.0030)	(0.0030)	(0.0030)	(0.0028)	(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 A13: US (NLSY) occupational decline and individual occupational stability

	(1)	(2)	(3)	(4)	(5)	(6)
A. Probability of working in same occupation in 2014 as in 1987 (mean: 0.09)						
Declining	-0.043	-0.039	-0.038	-0.032	-0.022	-0.012
	(0.017)	(0.017)	(0.017)	(0.017)	(0.017)	(0.019)
B. Probability of working in same occupation group in 2014 as in 1987 (mean: 0.20)						
Declining	-0.0050	0.00079	0.0025	-0.0087	0.016	0.022
	(0.032)	(0.028)	(0.029)	(0.031)	(0.029)	(0.026)
C. Probability of working in same broad occupation group in 2014 as in 1987 (mean: 0.36)						
Declining	-0.066	-0.048	-0.041	-0.044	0.033	0.041
	(0.052)	(0.042)	(0.043)	(0.044)	(0.027)	(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.