Job Polarization and Task-Biased Technological Change: Evidence from Sweden, 1975–2005*

Adrian Adermon[♣] and Magnus Gustavsson[♦]

Abstract

This paper shows that between 1975 and 2005, Sweden exhibited a pattern of job polarization with expansions of the highest and lowest paid jobs compared to middle-wage jobs. The most popular explanation for such a pattern is the hypothesis of 'task-biased technological change', where technological progress reduces the demand for routine middle-wage jobs but increases the demand for non-routine jobs located at the tails of the job-wage distribution. Our estimates, however, do not support this explanation for the 1970s and 1980s. Stronger evidence for task biased technological change, albeit not conclusive, is found for the 1990s and 2000s. In particular, there is both a statistically and economically significant growth of non-routine jobs and a decline of routine jobs. Results for wages are, however, mixed; while task-biased technological change cannot explain changes in between-occupation wage differentials, it does have considerable explanatory power for changes in within-occupation wage differentials.

Keywords: Inequality; Job Mobility; Skill Demand; Skill-Biased Technological Change

JEL Classifications: E24; J21; J23; J62; O33

* We are grateful to Per-Anders Edin, Maarten Goos, Nils Gottfries, Mikael Lindahl, Martin Nybom, three anonymous referees, and seminar participants at Uppsala University, the 2010 UCLS workshop, and the 2011 National Conference of Swedish Economists for useful comments. We wish to thank Maarten Goos, Alan Manning, and Anna Salomons for providing some of the data used in this paper. Magnus Gustavsson acknowledges financial support from the Swedish Council for Working Life and Social Research.

♣ Uppsala University, Department of Economics and UCLS, Box 513, SE-751 20 Uppsala, Sweden; Adrian.Adermon@nek.uu.se

♦ Uppsala University, Department of Economics and UCLS, Box 513, SE-751 20 Uppsala, Sweden; Magnus.Gustavsson@nek.uu.se

I. Introduction

Among economists, technological progress is commonly believed to increase labor demand for more-skilled workers relative to less-skilled workers. A major reason for this is the apparent fit between such ‘skill-biased technological change’ (SBTC) and the historical upward pressure on the returns to worker skills (for overviews, see Katz and Autor, 1999; Acemoglu and Autor, 2011). Less recognized is however that SBTC also has straightforward and important implications for the composition of jobs in an economy. In the typical textbook model, technology-induced shifts in labor demand that push the returns to skills above its long-run equilibrium will make it increasingly attractive for individuals to acquire skills—along the lines of standard human capital theory—and thus also produce a continuous increase in the supply of skills (e.g., Atkinson, 2008). Since there are increases in both the demand and supply of skills, ongoing SBTC predicts monotonic growth in the number of more-skilled to less-skilled jobs.

Recently, however, U.S., U.K., and German studies have documented a rising share of not only the highest-paid jobs but also of the lowest-paid jobs (e.g., Goos and Manning, 2007; Autor, Katz, and Kearney, 2008; Dustmann, Ludsteck, and Schönberg, 2009; Autor and Dorn, 2013). Assuming that wages can be thought of as a single-index of worker skills, this pattern is inconsistent with the implications of SBTC, where higher-paid jobs should simply increase relative to lower-paid jobs. Instead, as first demonstrated by Goos and Manning (2007), this pattern of ‘job polarization’—the disproportionate growth of *both* the lowest- and highest-paid jobs—is potentially more consistent with Autor, Levy, and Murnane’s (2003) (ALM henceforth) more nuanced version of technological change, which stresses the substitutability between routine tasks and technology and the notion of ‘task biased technological change’ (TBTC henceforth).

In their set-up, ALM make an important distinction between labor performing routine and non-routine tasks and argue that the falling price of computer power should yield a drop in the relative demand for labor performing routine tasks (e.g., bookkeepers, repetitive production work). This follows from the observation that computer-driven technology can primarily replace human labor in routine tasks—tasks that can be expressed by rules or step-by-step procedures—but not (as yet) in non-routine jobs. Goos and Manning (2007) in turn highlight that this fits well with job polarization since routine tasks are most common in middle-wage jobs. Top-paying jobs on the other hand consist of tasks that require non-routine cognitive skills (e.g., engineers, economists) which should be complementary to computers, and the bottom of the wage distribution consists of jobs with a high degree of non-routine manual tasks (cleaners, waiters, janitors) which, according to ALM, should be neither complements nor substitutes to computers. ALM's hypothesis combined with the observed job polarization thus implies a rise in the demand for low-wage workers relative to middle-wage workers and thereby—compared to traditional SBTC—offers a more nuanced view of how technology, and computers in particular, affects the demand for labor of different skills. In line with this, Firpo, Fortin, and Lemieux (2011) also conclude that changes in both between- and within-occupation wage differentials in the U.S. during the 1990s are in line with predictions from TBTC.

In light of these previous studies, the purpose of this paper is twofold. First, to thoroughly document the wage profile of net job creation in Sweden between 1975 and 2005, and second, to investigate if the observed job patterns are linked to the extent of routine versus non-routine tasks across the job distribution along the lines predicted by ALM's and Goos and Manning's (2007) TBTC hypothesis. In doing so, we provide three innovations to the empirical literature on TBTC. First, we use a bootstrap procedure to test if the observed pattern of net job creation is statistically significant. Tests of statistical significance are generally not carried out in previous studies, and our results show that such tests can affect conclusions. Second, we invoke longitudinal data to investigate if individual mobility across routine and non-routine jobs is along the lines expected from TBTC. Third, we apply Firpo, Fortin, and Lemieux's (2011) newly developed wage model to test if individuals' wage changes are in accord with TBTC, and as the first paper, base this test on longitudinal data. While Firpo, Fortin, and Lemieux (2011) use cross-sectional data for the U.S., they themselves and also Acemoglu and Autor (2011) recognize that longitudinal data is more likely to overcome econometric problems associated with workers' job mobility and self-selection into job tasks.

Since most of the previous research on job polarization pertains to the U.S. and U.K., Sweden is also a particularly interesting country to study because in many regards it lies at the opposite end of the institutional spectrum. In particular, Sweden has one of the world's most compressed wage structures, strong and influential unions, high levels of employment protection, and generous unemployment benefits combined with a well-developed welfare system (e.g., Cahuc and Zylberberg, 2004; Björklund and Freeman, 2010). Several studies have suggested that this could yield a different pattern of net job creation. Acemoglu (2001) shows within a matching framework of the labor market that generous unemployment benefits and high minimum wages—as can be found in Sweden—induce incentives that should shift the composition of employment towards high-wage jobs. On the other hand, Acemoglu (2003) suggests a model in which union-imposed wage compression encourages the adoption of technologies that increase the productivity of less-skilled workers and thus induces positive effects on labor demand for these groups. Acemoglu and Autor (2011) also discuss the possibility that powerful unions could restrict or delay the substitution of machines for tasks performed by labor. Hence, even though Sweden certainly could access the same technology as the U.S. and U.K., the marked differences in institutional preconditions need not imply job polarization in Sweden, even if TBTC is true for the U.S. and U.K.

Previewing the main results, we find that net job creation in Sweden does indeed display a pattern of job polarization over the full period 1975–2005. Dividing the analysis into the two sub-periods 1975–1990 and 1990–2005 does, however, show much stronger evidence for polarization in the later period. Our analysis of the relationship between routine and non-routine tasks across jobs and the observed changes in employment and wages also unable to provide statistical support for TBTC as an important explanation for the overall pattern of job creation in Sweden during the 1970s and 1980s. For the 1990s and 2000s, on the other hand, we do find significant relative declines of routine jobs and expansions of cognitive non-routine jobs both between and within industries—as would be expected if TBTC is a real phenomenon. Using the longitudinal dimension of our data, we also find a clear pattern of job mobility away from routine jobs towards cognitive non-routine jobs after 1990. Changes in within-occupation wage differentials after 1990 are also supportive of TBTC. However, we are unable to find support for TBTC when it comes to changes in between-occupation wage differentials

No previous study has made a formal statistical investigation of the connection between job tasks and the wage profile of employment creation in Sweden. In fact, most previous research, regardless of country, has primarily drawn conclusions based on visual inspections of distributions of routine and non-routine tasks across the wage ranking of jobs. Important exceptions are Goos, Manning, and Salomons (2009, 2010) who rely on a regression framework to investigate the cross-sectional connection between tasks and employment changes in Western Europe, as well as Kampelmann and Rycx (2011) who use regressions for Germany. Our corresponding estimates corroborate the findings of both these studies of a negative effect of routine tasks and a positive effect of cognitive non-routine tasks on job-specific employment during the 1990s.

Some previous studies have, to some extent, aimed at documenting the wage-quality of net job creation in Sweden—that is, the growth of “good” versus “bad” jobs. Fernández-Macías and Hurley (2008) use the European Union Labour Force Survey (ELFS) and report a pattern of skill upgrading—higher-paid jobs increase relative to lower-paid jobs—in Sweden since the mid 1990s. However, based on the same data source, Goos, Manning and Salomons (2009) instead report evidence of job polarization in Sweden over this period. A possible explanation for the contradicting results is differences in data processing.[1] Åberg (2004) uses Swedish data and finds a pattern of skill upgrading between 1977 and 2001. His applied sample is, however, small—the sample we use is more than twenty times larger—and we believe this to be the main explanation for the difference between his and our results.

Few previous studies have investigated the connection between job tasks and changes in occupation-specific wages. For the U.S., Firpo, Fortin, and Lemieux (2011) conclude that both changes in within- and between-occupation wage differentials during the 1990s are in accord with predictions from TBTC. For Germany, Kampelmann and Rycx (2011) investigate the connection between SBTC and between-occupation wage differentials since 1985, but are generally unable to find any statistically significant connection.

The rest of this paper proceeds as follows. Section 2 discusses the data and the empirical methodology. Section 3 first presents the wage profile of net job creation between 1975 and 2005 and then investigates its connection to routine versus non-routine tasks along the lines of TBTC. The paper ends with concluding remarks.

[1] The aim of these two studies is to provide broad overviews of occupational changes in a large set of European countries since the mid-1990s and this requires the data to be harmonized across countries. This harmonization differs across Fernández-Macías and Hurley (2008) and Goos, Manning and Salomons (2009). Their results for the rest of the investigated countries do also, to some extent, differ, with much stronger support for polarization across Europe in Goos, Manning and Salomons (2009).

II. Data and methodology

LINDA

The primary data for this paper comes from the Swedish longitudinal micro-database LINDA. Beginning in 1968, it contains a cross-representative sample of 3.3 percent of the Swedish population for each year (see Edin and Fredriksson, 1998, for details). We use data for three years: 1975, 1990, and 2005. Unlike for most other years, these three waves of LINDA contain detailed data on individuals' occupations, labor income, and hours worked. It is also possible to translate occupational classifications across these three years using official crosswalks developed by Statistics Sweden.[2]

LINDA is made up of different registers and surveys. For the years 1975 and 1990, we primarily use information collected from the Swedish Population and Housing Census ("Folk- och bostadsräkningen", FoB). For the year 2005, we primarily use information collected by Statistics Sweden through individuals' employers in the Linda Wage Survey. Individuals and employers are obligated by law to respond in their respective surveys. As a consequence, response rates are above 97 percent. An attractive feature of LINDA is its longitudinal dimension where, because of the link to registers and the very high response rates in the surveys, outflow occurs primarily because of death or migration from Sweden.

Net Job Creation and Job Polarization

Our approach to investigate net job creation in high-, middle-, and low-wage jobs builds on a methodology first proposed by Joseph Stiglitz while in the Clinton administration and later refined and extended by Wright and Dwyer (2003) and Goos and Manning (2007). In a first step, we define a job as a particular occupation in a particular industry. We use three-digit SSYK coding for occupation and two-letter SNI 2002 coding for industry. This gives an

[2] A written description of the translations as well as the used Stata do-files are available on request.

industry/occupation matrix with 3,503 job cells. Individuals in the age interval 18–64 years are placed in cells and weighted by their regular working time so that each cell contains the number of full time workers with a particular job. Since many cells are empty, we are left with 1,379 jobs for our analysis. These jobs contain nearly all individuals. The total sample sizes are 123,080, 124,120, and 117,535 individuals for 1975, 1990 and 2005, respectively.

In the next step, we rank jobs according to their median wage in the first year, 1975, and group them into quintiles based on their median wage and cell size in that year.[3] That is, we group jobs into the lowest paid 20 percent (quintile 1), the second lowest paid 20 percent, up to the top 20 percent based on their median wage and cell size in 1975.[4] To study net job creation in different parts of the wage distribution for jobs between 1975 and 2005, we compute changes in the number of jobs—individuals in a particular occupation in a particular industry—in each of the 1975 quintiles. In other words, the numbers of individuals in 1975 that have jobs that are in the lowest paid quintile are compared to the number of individuals in the same jobs in 2005.[5] This gives net job creation of the lowest paying jobs.[6] For instance, assume that 20,000 individuals (in full-time equivalents) are employed in the jobs that in 1975 were classified into the lowest paid quintile whereas in 2005 the same jobs hold 40,000 individuals (in full-time equivalents); this means that there has been a net job creation of the

[3] Experimentations with different ways to group jobs generally provided the same overall pattern of net job creation as our preferred choice of quintiles.
[4] Since we assign each job to a unique quintile, it is not possible to create groups that contain exactly 20 percent of the sample. Our "quintiles" thus contain between 19 and 21 percent of the sample each.
[5] In the data, some jobs disappear while new ones pop up in later years. Most such jobs have very few individuals in them and the great majority is due to statistical changes in how occupations are classified over time. In our main analysis, we only include those jobs that are present in 1975, but have performed several sensitivity analyses related to this. First, we have assigned jobs into quintiles based on their wage and employment in 2005 and then only included jobs that are present in 2005 (the opposite to our main approach). Second, we have only included jobs present in both 1975 and 2005. Third, we have as far as possible recoded (admittedly *ad hoc*) jobs that are new in 2005 into the 1975 classification. None of these approaches change our conclusions (results are available on request). In practice therefore, new and disappearing jobs do not seem to be a significant problem for our analysis.
[6] This interpretation is valid since the wage ranking of jobs is sufficiently stable over time. Between 1975 and 2005, the rank correlation for all jobs in our analysis is above 0.8, and assigning jobs to quintiles based on wages and employment in 2005 instead of in 1975 does not change any main results in our analysis. It is also worth noting that the relative composition of college educated workers across quintiles stays constant over time, indicating that the skill requirement, or skill ranking, in terms of formal education, stays constant over time; detailed numbers are available on request.

lowest paid jobs by 20,000 units. The same is done for jobs in each of the 1975 quintiles. In the analysis, we rescale job creation in our sample to match the aggregate changes for the whole of Sweden, so that the results can be interpreted as absolute growth for the whole of Sweden.[7]

As quantitative measures of job polarization, we use the percentage change in the ratio of employment in the first (lowest) job quintile relative to the third quintile, and in the fifth quintile relative to the third quintile. That is, taking the percentage change in employment at the first quintile (E^{q1}) relative to the third quintile (E^{q3}) between 1975 and 1990 as an example, we calculate the job polarization statistic as

$$(1) \qquad \%\Delta(E^{q1} / E^{q3}) = \frac{(E^{q1} / E^{q3})_{1990} - (E^{q1} / E^{q3})_{1975}}{(E^{q1} / E^{q3})_{1975}}.$$

This is also the statistic to which we apply our bootstrap.

The Bootstrap

Since our calculations rely on samples, it is desirable to not only present point estimates of job polarization but also on its statistical significance. The asymptotic distribution of the statistic in (1) is however unknown, but bootstrapping offers a means to approximate its finite sample distribution; see Horowitz (2001).

The bootstrap resampling is carried out in a way that suitably captures the temporal dependence in the data generating process for our LINDA-sample—that is, its longitudinal

[7] To translate changes in our sample into aggregate changes for the whole of Sweden, we use information on aggregate employment from Statistics Sweden. For each year, we first convert aggregate employment into full-time equivalents based on the distribution of hours worked in our LINDA sample. The number of individuals in our year-specific samples is thereafter rescaled to equal the aggregate number of full-time jobs in the economy for the same year. These rescaled samples are then used to calculate aggregate employment changes across quintiles.

dimension. We first pool the data for 1975, 1990 and 2005, resulting in a sample of n individuals. In each bootstrap replication, we with replacement draw a random sample of *individuals* of size n and keep all year-specific observations for each individual. This form of so-called 'block bootstrap' is motivated by the fact that the statistic of interest in (1) has a finite time dimension and asymptotically relies on $n \to \infty$; see Horowitz (2001) and Cameron and Trivedi (2005) for details.

We bootstrap the entire estimation procedure thus allowing for the stage-by-stage nature of our estimator. That is, in order to estimate (1), we first need to estimate the median wage in each job in 1975 (first-stage estimation), we next need to assign jobs into wage-employment quintiles in the 1975 distribution based on their median wages and employment in 1975 (second-stage estimation), and finally calculate (1) (third-stage estimation). So, for each bootstrap sample of n individuals, we perform all of these stages in the same way as with our original data.[8] Hence, the bootstrap takes account of the uncertainty associated with estimated median wages and the number of full time workers in each job in 1975, and thereby the thresholds used to divide jobs into quintiles, as well as the uncertainty associated with the employment changes in each quintile over time.[9]

Finally, for our bootstrap to be consistent, it must fulfill the so-called 'smoothness condition'; see Horowitz (2001). To analytically prove this condition in our application is far from straightforward. We have therefore instead checked this condition empirically by comparing our results to those produced by the "m out of n bootstrap", where the size of each bootstrap subsample m is less than the original sample size—that is, $m < n$. This 'subsample method' is less sensitive to violations of the smoothness assumption (Horowitz, 2001). In fact, Bickel, Götze, and van Zwet (1997, p.1) state that the m out of n bootstrap "has been known

[8] To be consistent with the construction of our original working sample, in each bootstrap and for all years we drop those jobs that are not present in the 1975 sample; see footnote 4.
[9] This way of dealing with stage-by-stage estimation is outlined in Wooldridge (2010) and Cameron and Trivedi (2005), and also applied in, for example, Dustmann and Meghir (2005).

to work in all known realistic examples of bootstrap failure". Cameron and Trivedi (2005, p.373) also state that "Subsample bootstraps are useful when full sample bootstraps are invalid, or as a way to verify that a full sample bootstrap is valid". As highlighted by Horowitz (2001), it is not a perfect substitute for the full bootstrap though, since it tends to be less accurate than the full bootstrap if the full bootstrap is indeed consistent.

The robustness of the *m* out of *n* bootstrap rests on the assumption that if $m \to \infty$, then $n \to \infty$ and $m/n \to 0$. Hence, as *n* grows, so should *m*, but at a sufficient slower rate. The literature does not yet offer any clear rule for the best choice of *m* in practical applications (e.g., Bickel and Sakov, 2008). We have therefore performed two set of *m* out of *n* bootstraps, one with *m* equal to 50 percent of our *n*, and another with *m* equal to 40 percent. These two sets of bootstraps produce confidence intervals that are very similar to each other and to the ones obtained by our full (main) bootstrap. We interpret this as empirical support for the consistency of our full bootstrap. An appendix containing the results of the *m* out of *n* bootstraps as well as a more detailed description of our bootstrap is available on request.

Task measures

To investigate the connection between routine and non-routine tasks and changes in employment—in light of the TBTC hypothesis—we use the three task measures developed and kindly provided to us by Goos, Manning, and Salomons (2009, 2010) of how intense occupations are in tasks labeled as *abstract*, *routine* and *service*. For our application, it should be noted these measures pertain to occupation classifications only and not occupation-industry combinations.[10]

[10] Occupations in O*NET are reported as 2000 Standard Occupation Codes (SOC), which Goos, Manning, and Salomon manually converted into the two-digit version of the International Standard for Classification of Occupations (ISCO-88). According to the authors, this translation was in all but a few cases straightforward since SOC is on a much more detailed level than ISCO-88; see Goos, Manning, and Salomons (2010) for details. Because of this, it is straightforward to merge these task measures to our data, since the two-digit version ISCO-88 is identical to two-digit SSYK.

The three task measures are constructed from 96 variables in the 2006-updated version of the US Occupational Information Network (O*NET) database. O*NET provides data on worker characteristics, worker requirements and general work activities for 812 U.S. occupations, information that in turn comes from job incumbents, occupational analysts and occupational experts.

Each of the 96 O*NET variables used by Goos, Manning, and Salomons contain information on the importance of a specific task across occupations on a 1–5 scale.[11] Each of these variables were categorised into one of the three categories *abstract*, *routine* or *service* based on the ALM-hypothesis of how well technology can substitute for the relevant task. For each of the three categories, the average of the included variables was then calculated for each occupation based on estimated principal components; see Goos, Manning, and Salomons (2010) for details. Selected descriptive statistics for the three task measures are presented in the appendix.

The link between computers and the three task measures are as follows. *Routine* tasks are intense in both cognitive and noncognitive routine skills and computers can perform these with relative ease, such as jobs that require the input of repetitive physical strength or motions, as well as jobs that require repetitive and non-complex cognitive skills. *Abstract* and *service* tasks are both in the non-routine dimension, but their skill content differs. *Abstract* tasks, such as "complex problem solving", are intense in non-routine cognitive skills and are expected to be complementary to computers. *Service* task, such as "caring for others", are intense in non-routine noncognitive skills and should not be directly affected by computerization. While *abstract* tasks are non-routine tasks mainly carried out by highly

[11] For some occupations, the task-content could potentially differ somewhat across the Swedish and U.S. labor markets. However, like Goos, Manning, and Salomons (2009, 2010) and Acemoglu and Autor (2001), we find it very unlikely that this difference should be large enough to disqualify the use of O*NET data in tests of TBTC for other OECD countries than the U.S.

educated workers (engineers and medical doctors), *service* tasks are non-routine tasks that workers with different levels of education may perform (medical doctors and hairdressers).

Examples of O*NET variables used as measures of *routine* tasks are the importance of "arm-hand steadiness", "manual dexterity", "operation monitoring", and "estimating the quantifiable characteristics of products, events or information". Examples of *abstract* task measures are "critical thinking", "judgment and decision making", "interacting with computers", and "thinking creatively". Examples of *service* task measures are "social perceptiveness", "service orientation", "selling", and "performing for or working directly with the public".[12]

III. Results

Net job creation across wage quintiles

Figure 1 displays net job creation across the five wage quintiles. There is a clear pattern of polarization with most of the employment growth occurring in the highest (fifth) and lowest (first) wage quintiles. In the appendix, we also display the most growing and declining jobs *within* each quintile between 1975 and 2005. In most quintiles, the most growing jobs are found in the education and health services, while the most shrinking jobs tend to be in agriculture, forestry and industrial production.

[12] Typical occupation groups with scores above average in *abstract* but below average in the other two measures include "Physicists, chemists and related professionals" and "Architects, engineers and related professionals". Occupation groups with scores above average only in *routine* include "Machine operators and assemblers" and "Labourers in mining, construction, manufacturing and transport". Occupation groups with scores above average only in *service* include "Personal and protective services workers" (e.g., police officers and cooks) and "Models, salespersons and demonstrators". Several occupation groups have above average scores on at least two of the task measures, including "Machinery mechanics and fitters" (*abstract* and *routine*), "Teaching professionals" (*abstract* and *service*), and "Drivers and mobile plant operators" (*routine* and *service*, e.g., taxi drivers). Task scores for all occupations in our data are available on request.

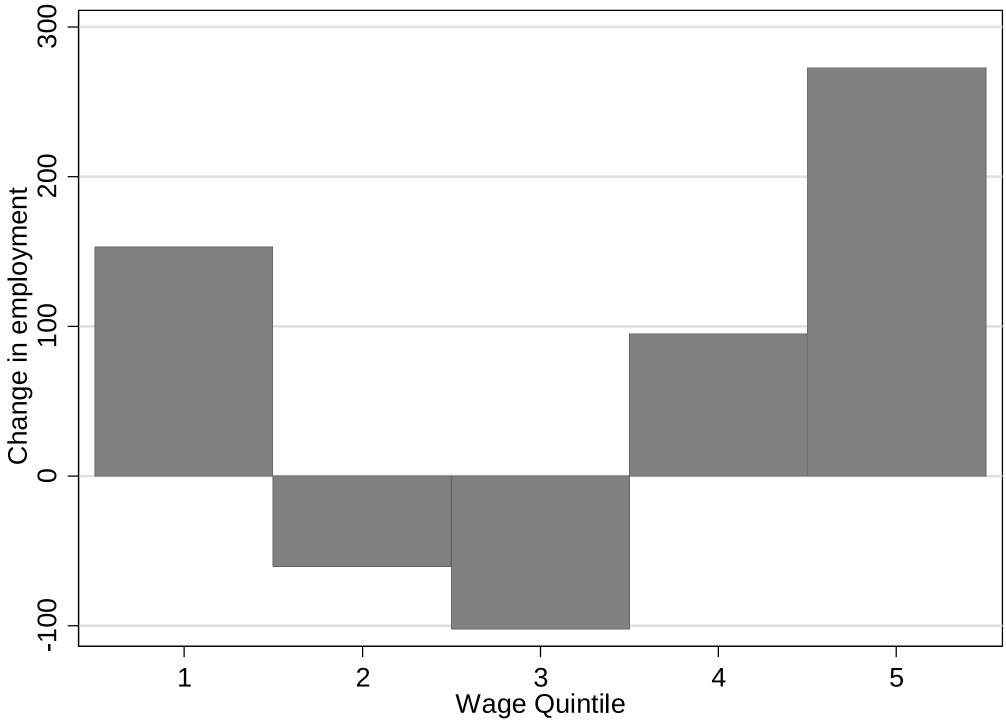

Figure 1: Change in employment by wage quintile, 1975–2005.

Note: The figure shows net job creation in each 1975 wage quintile. Changes in employment are in thousands of full-time equivalents, and are rescaled to match aggregate changes for the whole of Sweden. Quintile 1 is the lowest wage quintile.

Figures 2 and 3 further divide the changes into before and after 1990—the midyear in our sample. Both periods display polarization in the sense that employment in the middle quintile declines *relative* to jobs in the highest and lowest quintiles. On the aggregate level, the displayed changes fit well with previous knowledge about Swedish employment, with a steady growth of the employment to population ratio up until 1990, a sharp decline in connection with the severe economic crisis of the early 1990s followed by a rebound in the late 1990s but without reaching the pre-crisis level (e.g., Holmlund, 2006).

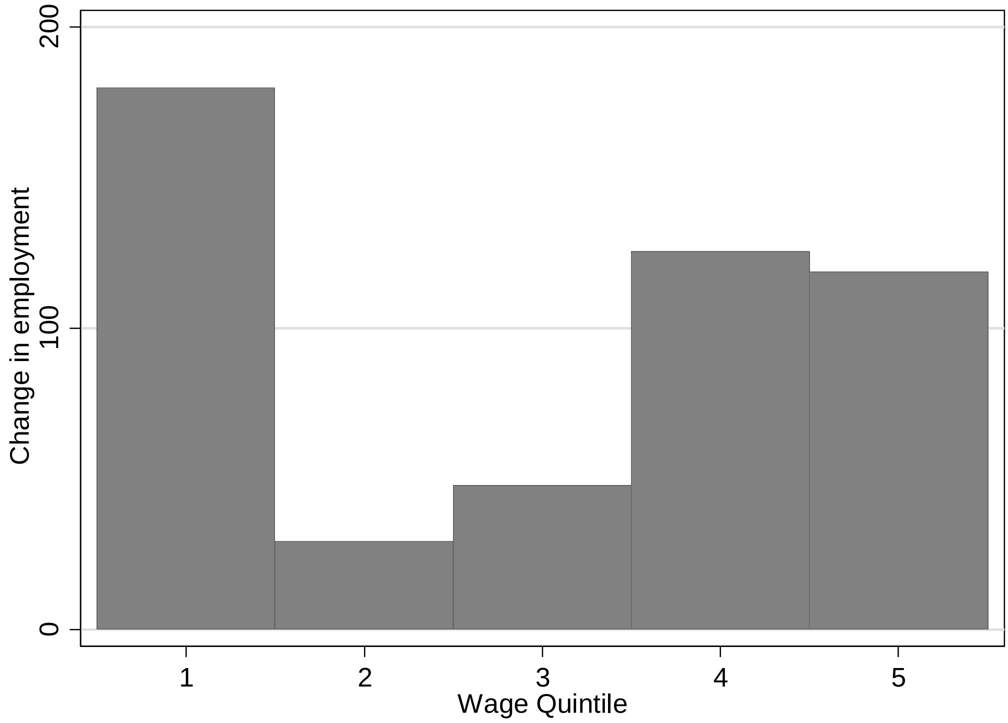

Figure 2: Change in employment by wage quintile, 1975–1990.

Note: The figure shows net job creation in each 1975 wage quintile. Changes in employment are in thousands of full-time equivalents, and are rescaled to match aggregate changes for the whole of Sweden. Quintile 1 is the lowest wage quintile.

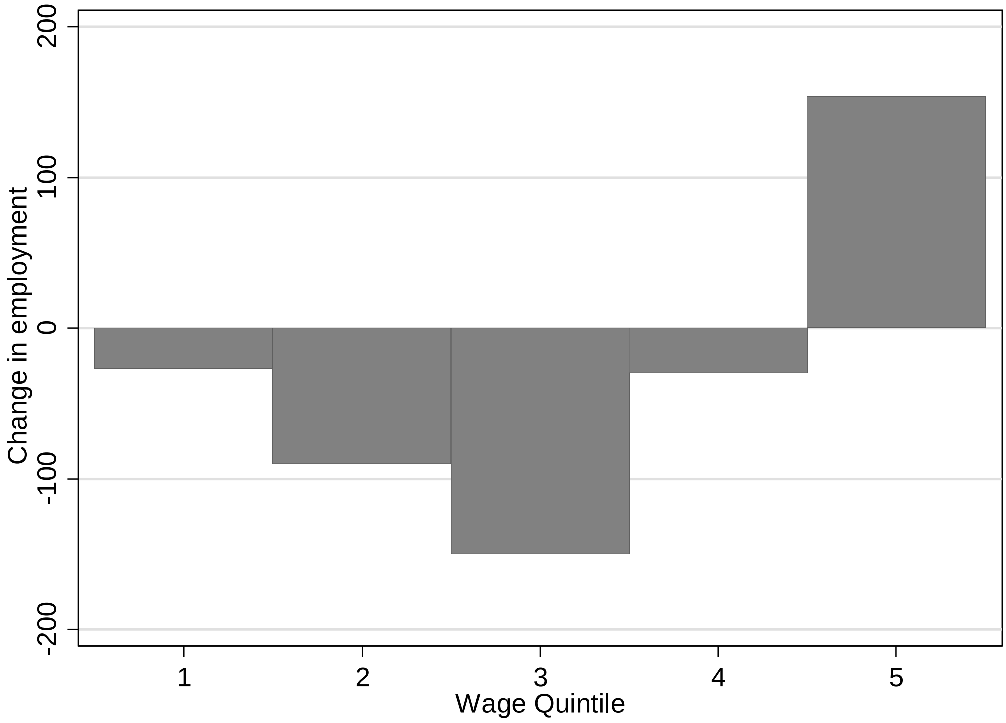

Figure 3: Change in employment by wage quintile, 1990–2005.

Note: The figure shows net job creation in each 1975 wage quintile. Changes in employment are in thousands of full-time equivalents, and are rescaled to match aggregate changes for the whole of Sweden. Quintile 1 is the lowest wage quintile.

Table 1: Economical and statistical significance of job polarization

	1975-2005	1975-1990	1990-2005
Quintiles			
$\%\Delta(E^{q1} / E^{q3})$	45.34	20.25	20.87
Bootstrapped 95% CI	[10.13, 64.66]	[-14.64, 32.66]	[10.30, 46.39]
$\%\Delta(E^{q5} / E^{q3})$	64.22	10.07	49.19
Bootstrapped 95% CI	[37.12, 88.54]	[0.482, 20.38]	[28.62, 64.48]
Tertiles			
$\%\Delta(E^{t1} / E^{t2})$	26.01	6.625	18.18
Bootstrapped 95% CI	[12.48, 40.84]	[0.444, 15.25]	[8.891, 27.27]
$\%\Delta(E^{t3} / E^{t2})$	53.68	11.98	37.24
Bootstrapped 95% CI	[35.28, 64.88]	[3.246, 18.05]	[26.81, 44.60]

Note: E^{q1} denotes employment in the lowest job quintile, and E^{t1} denotes employment in the lowest job tertile. Bootstrapped confidence intervals are in parentheses.

To clarify the extent of job polarization implied by Figures 1–3, the upper part of Table 1 displays the percentage changes in the ratio of employment in the lowest job quintile relative to the middle quintile, and in the highest quintile relative to the middle quintile; see equation (1). Between 1975 and 2005, jobs in the lowest quintile expanded by 45 percent relative to jobs in the middle quintile, with roughly equal contributions before and after 1990. The highest quintile expanded by 64 percent relative to the middle quintile over the same period, with most of the increase occurring after 1990.

Most previous studies in the literature on job polarization have relied on graphical analyses along the lines of Figures 1-3 without recognizing the statistical uncertainty of the estimated job pattern. Exceptions are Goos and Manning (2008) and Kampelmann and Rycx

(2011) who use *t*-statistics and regressions of percentile-specific changes in employment on linear and quadratic continuous percentile variables to test for polarization. We instead use the bootstrap methodology outlined in Section 2. A clear advantage with our approach is that it recognizes not only the uncertainty associated with changes in employment within a certain quintile but also the uncertainty associated with the initial division of jobs into these quintiles.

Bootstrapped 95 percent confidence intervals for the quintile ratios are contained in the upper part of Table 1. While all estimates for the periods 1975–2005 and 1990–2005 are statistically significantly different from zero, the null of a zero percentage change in (E^{q1} / E^{q3}) between 1975 and 1990 cannot be rejected. Thus, the statement above of an expansion of jobs in the lowest quintile relative to the middle quintile between 1975 and 1990 is associated with a noticeable amount of statistical uncertainty.

The bootstrapped confidence intervals might appear surprisingly wide. The main source for this is a great deal of uncertainty in the classification of some of the most dynamic jobs in the Swedish economy. For example, the most growing job in the whole Swedish economy over the period 1975–2005 is located in the first (lowest) quintile, but just below the cutoff for the second quintile.[13] In 40 percent of the bootstrap replications, this job is classified into the second instead of the first quintile (due to different estimates of median wages and the number of full time workers across replications). Consequently, these replications display markedly lower observed overall growth of the first quintile. Moreover, the most shrinking job in the third quintile is located just above the cutoff between the second and third quintile.[14] In 45 percent of the bootstrap replications, this job ends up in the second quintile instead of the third, yielding a more positive employment record for the third quintile in these replications.

[13] This is the job *Personal care and related workers* in *Health and social work*.
[14] This is the job *Managers of small enterprises* in *Wholesale and retail trade; repair of motor vehicles, motorcycles and personal and household goods.*

One way to reduce the statistical uncertainty in the estimates could be to divide jobs into wider percentile classes and thereby increase the sample size in each class. We therefore also present results based on dividing the job ranking into tertiles (thirds) instead and studying percentage changes in the lowest tertile (*t1*) relative to the middle tertile, and in the highest relative to the middle tertile. As can be seen in the lower half of Table 1, this approach reduces the statistical uncertainty and gives statistically significant estimates of polarization across the board. However, while the resulting job polarization based on tertiles over the full period 1975–2005 and for the sub-period 1990–2005 arguably is also economically significant, it is not obvious that this is the case for the period 1975–1990; the expansion of the lowest-paying jobs (first tertile) relative to middle-paying jobs is less than seven percent for the period 1975–1990.

A salient feature of the Swedish labor market is the high share and marked changes of public sector employment over time; there was a marked increase from 30 percent to over 40 percent of total employment during the 1970s followed by a decline to 35 percent during the 1990s. To investigate how this fits into the overall changes in the structure of employment, Figures 4 and 5 depict the patterns in the public and private sectors separately for the two periods 1975–1990 and 1990–2005, respectively.[15] In these figures, the classification of jobs into quintiles is based on the wage-employment ranking for the whole economy, so each job belongs to the same quintile as it did in Figures 1–3.

For the 1990s and 2000s, both sectors display a picture that more or less resembles the overall pattern during this period. Bootstrapped confidence intervals along the lines of those in Table 1 further confirm a statistically significant pattern of job polarization in both the public and private sectors (results are not shown but are available on request).

[15] The appendix contains a list of the ten largest jobs in the private and public sectors.

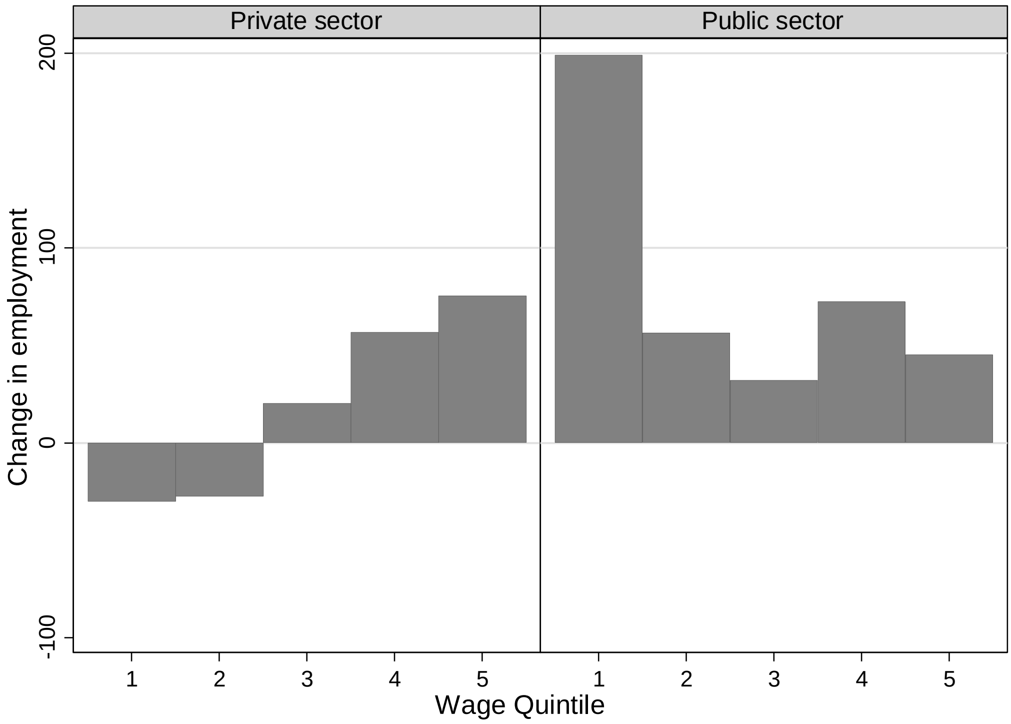

Figure 4: Change in employment by sector, 1975–1990

Note: The figure shows net job creation in each 1975 wage quintile by sector. Quintiles are calculated based on the full sample. Changes in employment are in thousands of full-time equivalents, and are rescaled to match aggregate changes for the whole of Sweden. Quintile 1 is the lowest wage quintile.

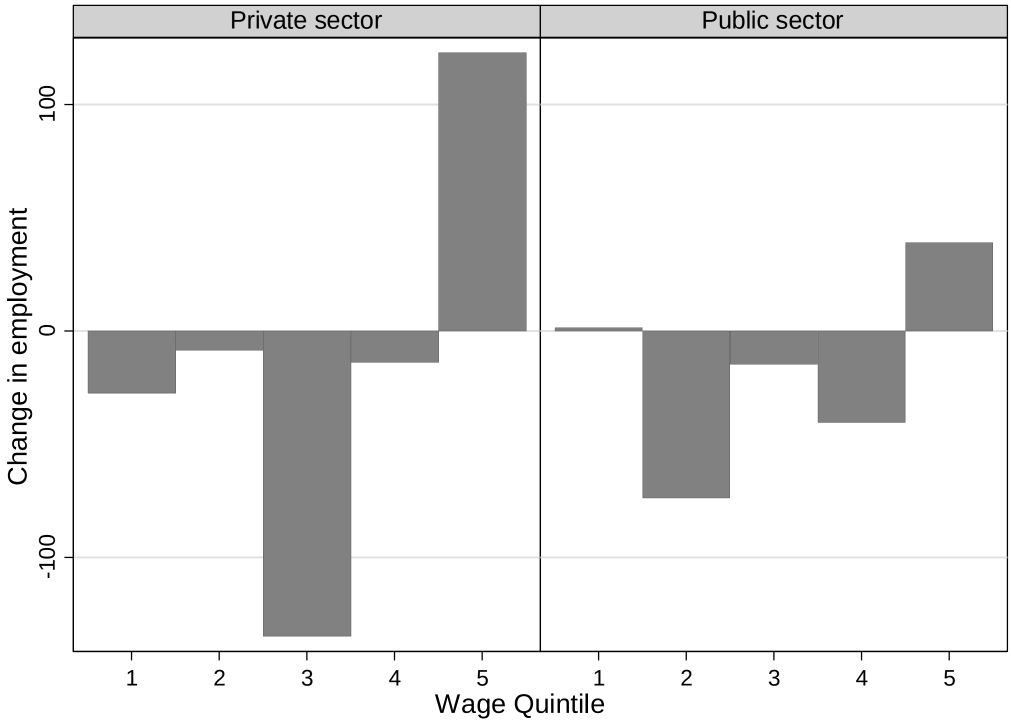

Figure 5: Change in employment by sector, 1990–2005

Note: The figure shows net job creation in each 1975 wage quintile by sector. Quintiles are calculated based on the full sample. Changes in employment are in thousands of full-time equivalents, and are rescaled to match aggregate changes for the whole of Sweden. Quintile 1 is the lowest wage quintile.

For the earlier period, 1975–1990, the patterns are markedly different across the two sectors though. The private sector displays much smaller changes and a pattern of skill upgrading with increases in higher paying jobs at the expense of lower paying jobs, and this skill upgrading is also statistically significant. The expansion of employment in the public sector does on the other hand appear to be largely driven by low-paying jobs.[16]

Another noticeable aspect of the Swedish labor market is the marked expansion of female employment during the 1970s and 1980s (e.g., Rosen, 1997). To investigate how this fits into the observed job polarization, Figures 6 and 7 present net job creation divided by gender (each job is again in the same quintile as in Figures 1–3). Between 1990 and 2005,

[16] A caveat with Figure 4 and the public sector, however, is the lack of any, based on bootstrapped confidence intervals, statistically significant changes.

both sexes contribute to the relative expansion of high- and low-wage jobs relative to middle-wage jobs. This pattern is highly statistically significant, using bootstrap confidence intervals. In comparison, between 1975 and 1990 hardly anything goes on with male employment whereas women display marked growth in all job quintiles. The lowest quintile does however account for a disproportional large share of the female employment expansion.[17]

Overall, the results for 1990–2005 show a pattern of job polarization in both the public and private sectors and for both female and male employment. This does not hold for the earlier period 1975–1990 though, as the public sector appears to have accounted for all of the overall growth in low-wage relative to middle-wage jobs. The division by gender further shows that these low-wage jobs were mainly filled by female workers. This is consistent with Rosen's (1997) findings that virtually all employment growth in Sweden during this period was accounted for by women entering employment into—the typically low paid—local government jobs. It is here also worth noting that according to the same author, these jobs were mainly the result of a political desire to expand the Swedish welfare state. If true, TBTC—or technological change in general—may have played a minor role in shaping the Swedish pattern of net job creation during the 1970s and 1980s. In the rest of the paper, we more directly investigate the explanatory power of TBTC by estimating the connection between tasks and changes in employment and wages across jobs.

[17] It should be noted, however, that the percentage change in the quotient between female employment in the first and third quintiles between 1975 and 1990 is not statistically significant. This is explained by our particular measure of polarization which is based on percentage changes. In 1975, the share of women in the first quintile is much higher than that in the third quintile, so even though there is a marked absolute increase in female employment in the first relative to the third quintile, the percentage increase in female employment in the first relative to the third quintile is much smaller.

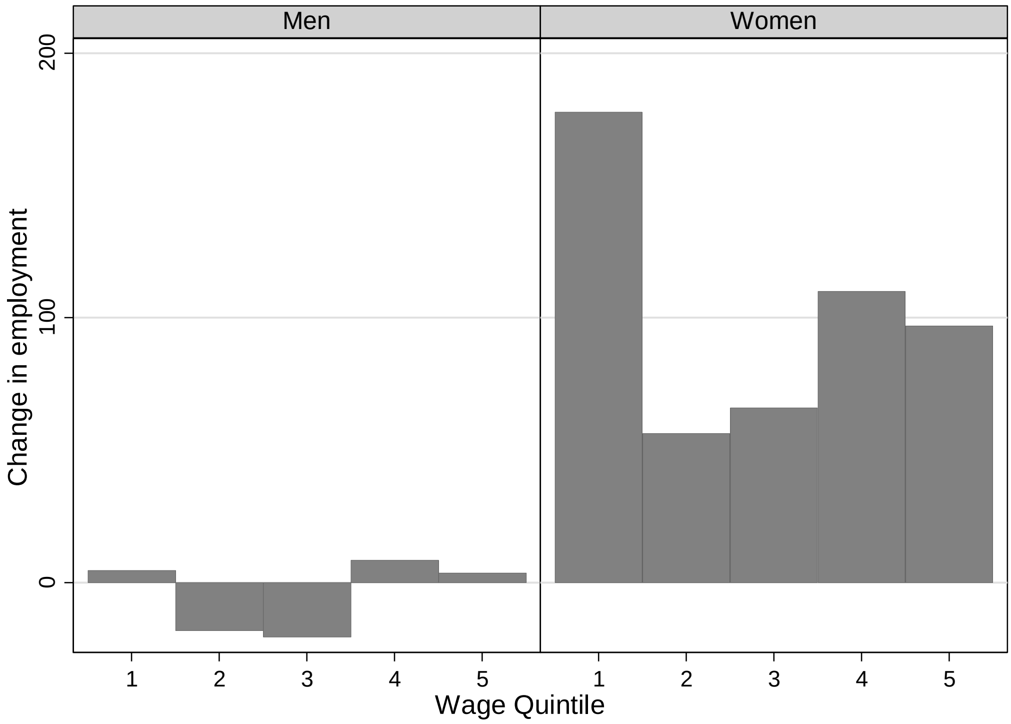

Figure 6: Change in employment by gender, 1975–1990

Note: The figure shows net job creation in each 1975 wage quintile by gender. Quintiles are calculated based on the full sample. Changes in employment are in thousands of full-time equivalents, and are rescaled to match aggregate changes for the whole of Sweden. Quintile 1 is the lowest wage quintile.

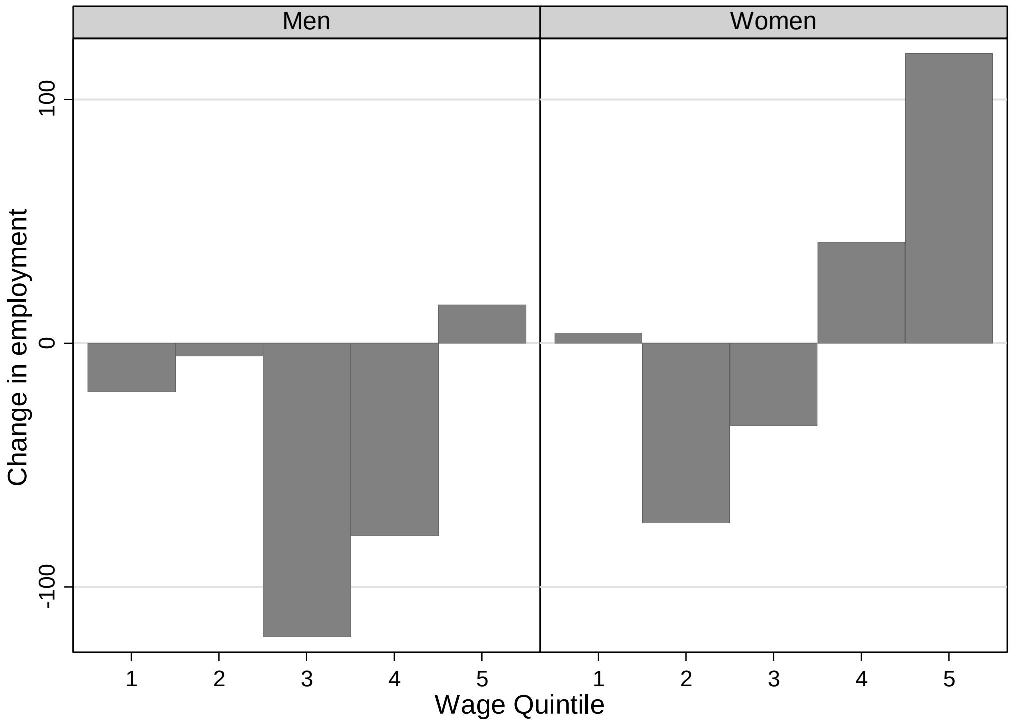

Figure 7: Change in employment by gender, 1990–2005

Note: The figure shows net job creation in each 1975 wage quintile by gender. Quintiles are calculated based on the full sample. Changes in employment are in thousands of full-time equivalents, and are rescaled to match aggregate changes for the whole of Sweden. Quintile 1 is the lowest wage quintile.

Job creation and routine versus non-routine tasks

As a first overview, Figure 8 displays the share of workers in each wage quintile in 1975 that are in an occupation with a task score on *abstract*, *routine* and *service* above the overall mean.[18] As can be seen, *abstract* tasks are more important in the highest paid jobs and *service* tasks are most important at the very highest and lowest paid jobs, whereas *routine* tasks are least important in the tails of the distribution. This mirrors previous documentations for other countries.[19]

[18] The overall mean for each task measure is calculated as the employment-weighted mean of that task measure across all jobs.
[19] We have also performed the analysis in Figure 8 using the five related routine and non-routine task measures developed by Autor, Levy and Murnane (2003), kindly provided by David Autor at http://econ-www.mit.edu/faculty/dautor. These are derived from much less information than those of Goos, Manning, and

We next regress changes in job-specific employment on the three task measures in the form of dummy variables that equal unity for occupations with scores above the overall mean.[20] To reduce sampling errors, we only perform regression based on jobs with at least ten employees in all three years 1975, 1990 and 2005; including all jobs in the regression analysis gives similar point estimates but substantially larger standard errors. Though we lose a lot of jobs by this restriction, we still retain 95 percent of all individuals in 1975 (93 percent in 1990, and 92 percent in 2005).[21]

The first five columns of Table 2 contain results for the period 1990–2005. The first column corroborates the impression from Figures 3 and 8 and is broadly consistent with TBTC; there is a statistically significant expansion of jobs intense in *abstract* tasks, a statistically significant decline in jobs intense in *routine* tasks, and no statistically significant change for jobs intense in *service* tasks.

Salomons (2009) but make it possible to investigate potential changes in job tasks over time since there are two versions of each measure, one created from information about job tasks in 1977 and one based on information from 1991. These alternative measures do not change any conclusions related to the distribution of routine versus non-routine tasks and they do not generally indicate marked changes in the extent of routine versus non-routine content of jobs over time; results are available on request.

[20] Using continuous task scores as explanatory variables produces the same conclusions in all analyses in this paper, both quantitatively and qualitatively (available on request). But since the continuous task scores have no intuitive meaning in themselves, we prefer the easy-to-interpret above-mean dummy variables.

[21] We do not weight the regressions by employment in the initial year. Since we use changes in employment as the dependent variable, we argue that also weighting by employment produces very hard-to-interpret results. For instance, suppose that most of the largest jobs in 1975 are intense in routine tasks, whereas most of the smallest jobs instead are intense in abstract tasks. If routine jobs then markedly decline to become the smallest jobs by 2005 whereas abstract jobs instead increase to become the largest jobs, then the weighting by initial employment will give a very large negative effect on the estimate for the dummy for routine tasks, but the positive estimate for abstract will be too small and most likely be statistically insignificant. One can easily think of other scenarios where weighting by employment gives strange results. While not using any weights has the drawback of not necessarily producing estimates that reflect average effects for the whole of Sweden, this is, as argued, probably also true for estimates that weight by initial employment. On the other hand, discarding weights will produce estimates that are straightforward to interpret, as they reflect the average effect on the log of employment across all jobs independent of initial employment.

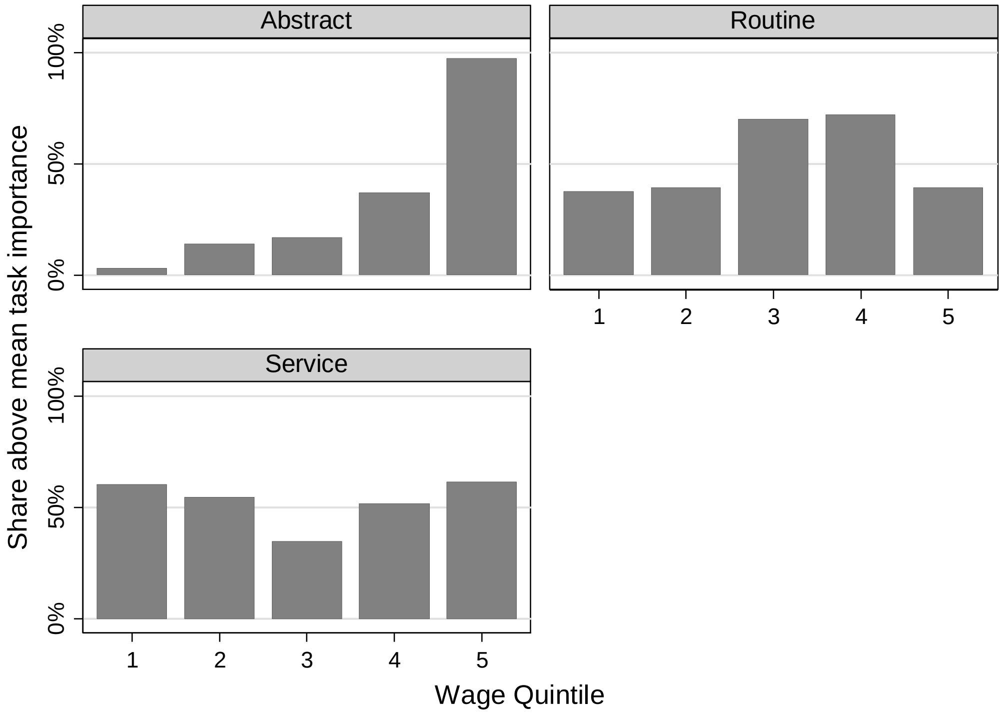

Figure 8: Incidence of abstract, routine, and service tasks across wage quintiles

Note: The figures display the share of workers in each job/wage quintile that are in an occupation with a task score above the overall mean. The underlying task scores are from Goos, Manning, and Salomons (2009)

In the next column, we add dummy variables for the 31 industries used to define a job (two letter level SNI) to try to control for changes in employment that stem from changes in product demand. That is, jobs that are intense in routine tasks may have declined in employment simply because these jobs are concentrated in industries with falling product demand and not necessarily because of organizational changes. The results show a statistically significant positive effect of *abstract* and a statistically significant negative effect of *routine* also *within* industries. Hence, industry-specific changes in employment are not behind the decline of routine jobs in Sweden after 1990.

The third column of Table 2 further adds the mean number of years of schooling (*education*) in each job in 1975 as a regressor. As argued by Goos, Manning and Salomons (2009, 2010), this variable allows for the predictions from the traditional SBTC hypothesis where employment should simply increase for jobs that require more education (more skills) relative to jobs that demand less education (less skills). As can be seen, *education* is not statistically significant and does not change any conclusions whereas a*bstract* is still significantly positive (at the 0.10 level) and *routine* is still significantly negative.[22]

[22] As stated above, this also holds if we instead use continuous task scores as regressors.

Table 2: OLS regressions; change in the log of job-specific employment

	1990-2005					1975-1990		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Abstract	0.282***	0.221***	0.154*	0.158*	0.143*	0.158**	0.137**	0.083
	(0.067)	(0.061)	(0.086)	(0.084)	(0.087)	(0.070)	(0.065)	(0.100)
Routine	-0.309***	-0.270***	-0.247***	-0.285***	-0.237***	-0.281***	-0.172**	-0.154*
	(0.084)	(0.080)	(0.081)	(0.082)	(0.081)	(0.091)	(0.087)	(0.086)
Service	0.012	-0.017	-0.030	-0.053	-0.023	-0.044	-0.061	-0.072
	(0.084)	(0.080)	(0.080)	(0.080)	(0.079)	(0.090)	(0.081)	(0.082)
Education			0.028	0.038	0.028			0.022
			(0.025)	(0.025)	(0.025)			(0.031)
Highly offshorable				-0.266**				
				(0.123)				
Offshorable					0.093			
					(0.063)			
Industry		X	X	X	X		X	X
Observations	478	478	478	478	478	478	478	478
R^2	0.086	0.260	0.262	0.269	0.265	0.041	0.237	0.238

Note: Dependent variable is change in log job-specific employment. Robust standard errors are in parentheses. Regressors are dummy variables equal to one if the skill measure for a job is above the mean. Education is mean education in a job in 1975.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

In the U.S., it has been argued that the decline of routine jobs in the middle of the wage distribution could be due to increased offshoring of jobs, that is, the migration of employment from the home country to other (mostly poorer) countries, rather than due to substitution between labor and computers along the lines of TBTC (see the discussion in Acemoglu and Autor, 2011). There is a lack of data on the number of jobs actually offshored for most countries, but recent attempts to classify the *offshorability* of jobs—the feasibility to perform the work duties from abroad—by Blinder (2009) and Blinder and Krueger (2013) actually suggest that there is no correlation between the extent of routine tasks in a job and its offshorability. For instance, Blinder and Krueger (2013, p.S127) conclude that "routine work is no more likely to be offshorable than other work".

To see if the offshorability of jobs could still potentially change any of our conclusions for the period 1990–2005, we use Blinder's (2009) classification of a job's offshorability, which in turn is based on information in the U.S. O*NET database. Blinder (2009) categorizes jobs into one of four levels of offshorability: 1) *Highly offshorable* (a person in this job does not have to be physically close to a work unit, e.g., computer programmers and telemarketers); 2) *Offshorable* (the whole work unit could be moved abroad, e.g. most factory workers); 3) *Non-offshorable* (whole work unit must be in home-country, e.g., sales managers), 4) *Highly non-offshorable* (e.g., child-care workers and farmers). The reader is referred to Blinder's study for more information on the criteria underlying this classification.[23]

In the fourth column of Table 2 we include, for the period 1990–2005, a dummy variable for jobs that are judged to be *highly offshorable*. This variable is negative and statistically significant, but its inclusion does not change the results for the other variables. In the last column for 1990–2005, we instead include a dummy for all jobs that are *offshorable*

[23] Correlations between our measures of jobs' tasks and offshorability are presented in the appendix.

(which includes *highly offshorable* jobs). This variable is insignificant, and other estimates are unchanged. Based on this, combined with previous studies in the literature, we find it unlikely that offshoring can explain the polarization of the Swedish labor market between 1990 and 2005. Overall, our set of regressions for the period 1990–2005 are broadly in line with the results for Western Europe in Goos, Manning, and Salomons (2009, 2010).

We next present regression results for the period 1975–1990. The last three columns of Table 2 present specifications with, first, the three task measures as the only regressors, then with industry dummies added, and finally also with the inclusion of average educational attainment in 1975. As expected, the evidence for TBTC is weaker compared to the period 1990–2005, although not completely absent. The estimates for *abstract* and *routine* are statistically significant and with the expected sign in the first two regressions, whereas only *routine* is statistically significant (at the 0.10 level) once *education* is controlled for. The share of the total variation explained by the three task measures is lower for this period; the R^2 with only the three task dummies is less than half of that for 1990–2005 and the partial R^2 for these three task measures when industry dummies are included is also three times smaller for the period 1975–1990, with a value of 0.06 for 1990–2005 versus 0.02 for 1975–1990.[24]

We have also included the variables *offshorable* and *highly offshorable* for the period 1975–1990 (not shown). Their estimates are positive and statistically significant, without affecting the estimates for the other variables. A possible explanation for the positive estimates is the fact that Blinder's (2009) measures are created with modern information technology in mind, and thus may be a bad proxy for the offshorability of jobs prior to the 1990s.

We have also performed the same analysis as in Table 2 for the private and public sectors separately. Estimates for the private sector are supportive of TBTC after but not prior

[24] This is the same kind of partial R^2 used in Acemoglu and Autor (2011). A formula and explanation can be found in Kennedy (1998).

to 1990. Estimates for the public sector are imprecise, but the sign of the point estimates are generally in line with those for the whole economy in Table 2. Overall, the results do not alter any of the conclusions regarding TBTC from Table 2.

A caveat with interpreting the results in Table 2 as evidence against or in favor of TBTC is that the estimates could be driven by exogenous changes in the composition of the labor force—labor supply—rather than by changes in relative labor demand. For instance, the large inflow of refugee immigrants into Sweden could affect the job composition, and the results in the previous sub-section also suggest the rise in female labor force participation as an important explanation for the observed job polarization up to 1990.

Our access to longitudinal data here allows us to shed light on the importance of changes in labor force composition, since we can investigate if individual job mobility is also in line with the results in Table 2. To do this, we re-estimate the regressions in Table 2, but only base the calculation of the dependent variable on those individuals that held a job in both the start and end of the investigated periods. Changes in job-specific employment—the dependent variable—can hence only be driven by differences in the extent of job mobility to and from different kinds of jobs, and not by entries and exits from the labor market/employment. That is, changes in the composition of the labor force will not affect these estimates.

The longitudinal estimates will not be without potential flaws though. They will probably not provide statistical evidence in favor of TBTC if technological change primarily affects the type of jobs available to new entrants on the labor market rather than incumbent workers, or if the main effect of TBTC is to increase non-employment among former routine workers. An additional caveat with our longitudinal approach is that all included individuals become 15 years older over the two periods 1975–1990 and 1990–2005, so the observed job movements could potentially reflect some sort of “career effect”. With this in mind, the longitudinal and cross-sectional estimates should be viewed as complements, as they have different strengths and weaknesses.

Table 3 presents the longitudinal estimates. For the period 1990–2005, the longitudinal estimates are similar to the cross-sectional estimates in Table 2 and imply statistically significant individual job mobility away from jobs with *routine* tasks towards jobs with *abstract* tasks, even when mobility between industries are accounted for. This constitutes additional evidence in favor of TBTC during this period. This conclusion, however, does not carry over to the estimates for 1975–1990. In particular, the estimate for *routine* is not statistically significant when we control for industry-specific effects. We are hence not able to reject the hypothesis that the expansion of certain industries, rather than organizational changes within industries, can account for mobility away from routine jobs during this period.

Table 3: OLS regressions; longitudinal changes in log job-specific employment

	1990-2005					1975-1990		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Abstract	0.226***	0.157**	0.196**	0.212**	0.198**	0.180**	0.144**	0.168*
	(0.070)	(0.067)	(0.093)	(0.091)	(0.094)	(0.070)	(0.067)	(0.100)
Routine	-0.217***	-0.218***	-0.232***	-0.285***	-0.234***	-0.147*	-0.068	-0.077
	(0.082)	(0.082)	(0.085)	(0.087)	(0.086)	(0.089)	(0.090)	(0.093)
Service	0.077	0.031	0.037	0.010	0.035	-0.019	-0.046	-0.042
	(0.082)	(0.080)	(0.081)	(0.081)	(0.081)	(0.088)	(0.082)	(0.083)
Education			-0.017	-0.007	-0.017			-0.010
			(0.026)	(0.026)	(0.026)			(0.030)
Highly offshorable				-0.359***				
				(0.120)				
Offshorable					-0.016			
					(0.068)			
Industry		X	X	X	X		X	X
Observations	478	478	478	478	478	478	478	478
R^2	0.057	0.192	0.193	0.206	0.193	0.023	0.177	0.177

Note: Dependent variable is change in log job-specific employment for individuals with employment in both 1975 and 1990, and in both 1990 and 2005. Robust standard errors are in parentheses. Regressors are dummy variables equal to one if the skill measure for a job is above the mean. Education is mean education in a job in 1975.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

How do the correlations between tasks and employment in Tables 2 and 3 fit with the observed job polarization between 1990 and 2005, as displayed in Table 1? To try to investigate this, we use the regression estimates to see how important different tasks across the job distribution are for the observed polarization. We first use the estimates in Table 2 to calculate the degree of polarization that would prevail if all jobs would have had zeros on the dummy variables *abstract*, *routine*, and *service*. We then compare this to the actual amount of polarization. The difference between the actual and counterfactual polarization is then used to calculate the explanatory power of tasks for the observed polarization.

As our measures of polarization, we use the percentage changes in the ratio of employment in the first quintile relative to the third quintile, (E^{q1}/E^{q3}), and in the fifth quintile relative to the third quintile—the same measures outlined in equation (1) and used in Table 1. To calculate the counterfactual extent of polarization that would remain if all jobs had zeros on the three dummy variables, we first convert the estimates in Table 2 for *abstract* and *routine* into the implied percentage effects (the anti-log of the estimate minus unity); we do not use the estimate for *service* since it is statistically insignificant. Denote these percentage effects for *abstract* and *routine* by γ_a and γ_r, respectively. To obtain a counterfactual value for employment in the first quintile in 2005, denoted $\tilde{E}_{05}^{q1}$, we subtract the change in employment since 1990 associated with the extent of *routine* and *abstract* tasks in the first quintile from the actual value of E_{05}^{q1}, using the formula

$$\text{(2)} \qquad \tilde{E}_{05}^{q1} = E_{05}^{q1} - E_{90}^{q1}(\gamma_r \cdot \overline{routine^{q1}} + \gamma_a \cdot \overline{abstract^{q1}}),$$

where $\overline{routine^{q1}}$ is the mean value of *routine* in the first quintile, that is, the share of employment in the first quintile with the dummy variable *routine* equal to unity. In equation

(2), the positive change in employment between 1990 and 2005 associated with *abstract* tasks in the first quintile is removed, as is the negative effect associated with *routine* tasks. What is left is the level of employment in the first quintile in 2005 that would exist if all jobs had the same task content (as measured by the dummy variables *abstract* and *service*). The corresponding calculation is done for the other quintiles. To obtain a counterfactual value of polarization between 1990 and 2005, we then use the resulting values of $\tilde{E}_{05}^{q1}$ and $\tilde{E}_{05}^{q3}$ in the calculation of percentage changes in the ratio of employment in the first job quintile relative to the third quintile, calculated as $[(\tilde{E}_{05}^{q1}/\tilde{E}_{05}^{q3})-(E_{90}^{q1}/E_{90}^{q3})]/(E_{90}^{q1}/E_{90}^{q3})$, and so forth for the earlier period and for changes in employment in the fifth relative to the third quintile.

Between 1990 and 2005, the actual increase of employment in the fifth relative to the third quintile was 49 percent—that is, (E^{q5}/E^{q3}) increased by 49 percent. Based on the estimates for *abstract* and *routine* in the specification with only the three task dummies in the first column of Table 2 and the formula in equation (2), the counterfactual increase in (E^{q5}/E^{q3}) is 5.3 percent. That is, when we replace the actual values of E_{05}^{q3} and E_{05}^{q5} by their counterfactual values $\tilde{E}_{05}^{q3}$ and $\tilde{E}_{05}^{q5}$ that are cleansed of the impact of job tasks, we observe a much smaller growth of the fifth relative to the third quintile. In fact, the obtained number implies that the distribution of tasks can potentially account for 89 percent (1-5.3/49) of the actual percentage increase in (E^{q5}/E^{q3}); see Table 4.[25] When we instead use the estimates for *abstract* and *routine* from the specification with industry-specific effects in the second column of Table 2, we get a counterfactual increase in (E^{q5}/E^{q3}) of 21.5 percent. This means that the share of the actual increase explained by job tasks decreases to 56 percent (1-21.5/49) once we take industry effects into account. For the period 1975–1990, we obtain

[25] The explanatory power of 89 percent may seem strange since R^2 in the underlying regression implies an explanatory power of 8.6 percent. However, R^2 measures the share of variance explained both within and between the job quintiles, whereas our measure only applies to the variation between quintiles.

negative counterfactual percentage changes in (E^{q5}/E^{q3}), so tasks can potentially account for all of the—moderate—actual increase in (E^{q5}/E^{q3}) between 1975 and 1990.

We next turn to the explanatory power of tasks for changes in (E^{q1}/E^{q3}), which we view as the most interesting exercise since the stand-out prediction of the TBTC hypothesis—compared to that of traditional SBTC—is the expansion of the lowest-paid jobs relative to middle-paid jobs. Based on the specification with only the three task dummies in the first column of Table 2, for the period 1990–2005 we find, using the same methodology as above, that the distribution of tasks can account for 36 percent of the actual increase in (E^{q1}/E^{q3}). Using estimates from the specification with industry-specific effects in the second column raises the share explained further, to 44 percent. For the period 1975–1990, the specification with only the three tasks measures can potentially account for 35 percent of the (statistically insignificant) actual increase in (E^{q1}/E^{q3}). Taking industry effects into account does however markedly lower the explanatory power to 14 percent for the period 1975–1990.

Table 4: Share of relative employment change explained by task content

	$\Delta(E^{q1}/E^{q3})$		$\Delta(E^{q5}/E^{q3})$	
	90–05	75–90	90–05	75–90
Actual change (%)	21	20	49	10
Share of change explained (%)				
Without industry dummies	36	35	89	100
With industry dummies	44	14	56	100

Note: The share of change explained is calculated based on regression estimates in Table 2; see the text for details.

TBTC and wages

According to the results in the previous sub-section, correlations between job tasks and employment changes are consistent with the idea of TBTC as an important phenomenon in Sweden during the 1990s and 2000s, but probably not during the 1970s and 1980s. As such, one expects technology-induced changes in the demand for labor along the lines of TBTC to also affect relative wages after 1990.

The major changes in the Swedish wage structure during the last 40 years are as follows. There was precipitous wage compression in all parts of the distribution from the late 1960s through the 1970s followed by a slight trend increase during the second half of the 1980s. These changes correspond well to the rise and fall of the union-induced egalitarian wage policy—'the solidarity wage policy'—though there also appear to be some room for market-driven explanations (Edin and Holmlund, 1995; Edin and Topel, 1997). After the 1980s, there is a clear upward trend in Swedish wage dispersion with most of the increase occurring in the upper half of the distribution (Gustavsson, 2006; Domeij, 2008).

Unfortunately, the effect of TBTC on relative wages is not straightforward to investigate, especially in contrast to the hypothesis of SBTC. With classical SBTC, there is no role for tasks, only individual skills, such as years of schooling, and this allows for a division of individuals into clear skill categories, such as high and low skilled labor. Such a division, in turn, makes it possible to use an empirical 'Katz-Murphy CES-model' (Katz and Murphy, 1992) to investigate the effect of demand shifts on relative wages, holding labor supply constant. Once different tasks are introduced, and individuals differ in their ability to perform these tasks, it is however no longer straightforward to simply divide individuals into skill groups (see Acemoglu and Autor, 2011, for a discussion of this). On the other hand, the Katz-Murphy model is largely unable to satisfactorily account for changes in the U.S. wage dispersion after the 1980s, which is one reason for the recent interest in TBTC (Acemoglu and

Autor, 2011). Because of the empirical difficulties associated with TBTC and wages, we in this sub-section use two somewhat different strategies to test if changes are consistent with demand shifts along the line of TBTC, each method with its own strengths and weaknesses.

First, we re-estimate the specifications in Table 2 but with changes in the log of median wages as the dependent variable instead of changes in the log of employment. The results are presented in Table 5.[26] For the period 1990–2005, results for the three task measures are neither economically nor statistically significant. Only the dummy *Highly offshorable* is statistically significant for this period, but with the "wrong" sign. In short, these results do not provide any statistically significant evidence in favor of TBTC as an important factor in the Swedish wage setting during the 1990s and 2000s.[27] As can be seen in Table 5, this also holds for the period 1975–1990. Moreover, performing the analysis in Table 5 for the private and public sectors separately does not alter any conclusions regarding TBTC (results are not shown but are available on request).

[26] We do not include any controls for changes in the composition of observable individual characteristics across jobs. Under the assumption that TBTC is true, such controls are likely to be outcome variables of the three task measures. For instance, changes in the share of immigrants in a job is an outcome variable if TBTC makes it easier for new entrants on the Swedish labor market to find employment in *service* or a*bstrac*t jobs rather than in *routine* jobs. Including an outcome variable as a regressor will bias the results for the task measures, and constitute what Angrist and Pischke (2009) refer to as a 'bad control'.

[27] Using changes in mean wages rather than median wages does not change this conclusion.

Table 5: OLS regressions; change in the log of job-specific median wages

	1990-2005					1975-1990		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Abstract	0.006	0.018	0.020	0.019	0.019	-0.045***	-0.044***	-0.021
	(0.012)	(0.011)	(0.015)	(0.015)	(0.015)	(0.013)	(0.013)	(0.017)
Routine	0.000	-0.008	-0.009	-0.001	-0.008	0.008	0.006	-0.002
	(0.012)	(0.012)	(0.014)	(0.014)	(0.014)	(0.021)	(0.019)	(0.019)
Service	0.007	0.011	0.011	0.016	0.012	-0.016	-0.022	-0.017
	(0.012)	(0.012)	(0.012)	(0.012)	(0.012)	(0.021)	(0.019)	(0.020)
Education			-0.001	-0.003	-0.001			-0.010*
			(0.006)	(0.006)	(0.006)			(0.006)
Highly offshorable				0.057**				
				(0.024)				
Offshorable					0.010			
					(0.011)			
Industry		X	X	X	X		X	X
Observations	478	478	478	478	478	478	478	478
R^2	0.001	0.157	0.157	0.170	0.159	0.035	0.084	0.091

Note: Dependent variable is change in log job-specific median wage. Robust standard errors are in parentheses. Regressors are dummy variables equal to one if the skill measure for a job is above the mean. Education is mean education in a job in 1975.

* p<0.10, ** p<0.05, *** p<0.01

However, as highlighted by Acemoglu and Autor (2011), regressions like those in Table 5 are, if TBTC is indeed true, potentially biased in a direction that rejects TBTC. If, as a result of TBTC, individuals who have a comparative advantage in routine tasks are laid off from such jobs and forced to move to jobs that rely more heavily on abstract or service tasks, their wages are likely to fall. Because of this, average or median wages associated with abstract and service tasks may also fall. Hence, cross-sectional estimates for *abstract* and *service* could be downward biased.[28] Estimated effects of *routine* on median wages will on the other hand be biased upwards if the least skilled workers, with lowest wages, are the ones that are laid off.

To try to overcome these statistical problems, we employ the empirical approach suggested by Firpo, Fortin, and Lemieux (2011). Based on a Roy-type linear skill-pricing model, they demonstrate how SBTC should affect both between- and within-job differentials and how this could be tested with longitudinal data. Of course, this approach is itself not without drawbacks; we discussed this further below.

The starting point in Firpo, Fortin, and Lemieux's (2011) model is a clear distinction between skills and tasks. Each worker comes with a bundle of skills to be used in a single occupation, consisting of a set of tasks. Since these skills cannot be unbundled and efficiently allocated across occupations, wages are likely to vary across occupations *conditional* on the skills of workers, as in a standard Roy model. As an example, consider two occupations, mathematicians and movers, and two skills, cognitive and physical strength. Obviously, physical strength is essential for movers whereas cognitive skills have a large impact on the productivity of mathematicians. On the other hand, physical strength is not important for the productivity of mathematicians, if it matters at all. Although people who choose to be

[28] Note that a standard longitudinal fixed-effects regression cannot overcome this problem, since the price of the fixed effect is not constant across occupations, and probably neither over time if TBTC is indeed true. For completeness and to get descriptive statistics, we have nevertheless, for the two periods 1975–1990 and 1990–2005, estimated the very basic panel regression $\Delta \ln w_{it} = \alpha + \beta_1 \Delta Abstract_t + \beta_2 \Delta Routine_t + \beta_2 \Delta Service_t + e_{it}$. The estimate for *abstract* is positive and significant for both periods, *routine* is negative and significant for the later period, and *service* is positive and significant for the earlier period; the other estimates are statistically insignificant. Detailed results are available on request.

mathematicians generally have a high ratio of cognitive to physical skills whereas the opposite is true for movers, Firpo, Fortin, and Lemieux (2011) forcefully argue that individuals are not heterogeneous to the point that the marginal product of these two skills are equalized across occupations. Instead, there will be an oversupply of physical strength among mathematicians, driving the returns to physical strength to zero among them. Likewise, there will be an oversupply of cognitive skills among movers that will drive the return for this skill toward zero.

Now, since the returns to physical skills for mathematicians is already driven to zero, technological change that allows for the substitution of human strength for cheap machines will not affect wages for mathematicians. For movers, on the other hand, such technological change should have a large effect on wages. The same reasoning applies to technological change that affects the productivity of cognitive skills: it should have a large impact on wages among mathematicians but no or a very minor effect on wages among those working as movers. In essence, such heterogeneous effects are used to identify changes in the demand for *routine*, *abstract* and *service* skills in Firpo, Fortin, and Lemieux's (2011) empirical model.

In detail, the log wage for worker *i* in occupation *j* at time *t* is assumed to be set according to

(3) $$w_{ijt} = \theta_{jt} + \sum_{k=1}^{K} r_{jkt} S_{ik} + u_{ijt} \,,$$

where S_{ik} (for $k = 1,...,K$) is each skill component k embodied in worker *i* and r_{jkt} are the occupation-specific returns to each skill component, θ_{jt} is a base payment that a worker receives in occupation *j* regardless of her skills, and u_{ijt} is an idiosyncratic error term. Since no individual-level data on the S_{ik} exists, changes in the returns to skills over time cannot be directly estimated. Instead, we use our longitudinal data and individuals who stay in the same occupation *j* over time and estimate

$$\text{(4)} \qquad \Delta w_{ij} = a_j + b_j w_{ij0} + e_{ij},$$

where w_{ij0} is individual i's wage in a base period ($t = 1975$ for changes between 1975 and 1990 and $t = 1990$ for changes between 1990 and 2005). Under the simplifying assumption that the different skill components S_{ik} are uncorrelated, Firpo, Fortin, and Lemieux (2011) show that

$$\text{(5)} \qquad b_j = \frac{\text{cov}(\Delta w_{ij}, w_{ij0})}{\text{var}(w_{ij0})} = \frac{\sum_{k=1}^{K} (r_{jk0} \Delta r_{jk}) \sigma_{kj}^2}{\sum_{k=1}^{K} r_{jk0} \sigma_{kj}^2 + \sigma_{uj0}^2}.$$

Equation (5) implies that we are able to learn something about changes in the r_{jkt}'s from the estimated slope coefficient in (4). While the denominator in (5) is always positive, the sign of the numerator depends on the correlation between returns to skills in the base period (r_{jk0}) and changes in the return to skill (Δr_{jk}). Based on the same kind of arguments as above concerning returns to skills for individuals working as movers versus mathematicians, TBTC implies that b_j should be negative for occupations where *routine* skills are very important. That is, in occupations where the return to workers' "routine skills" used to be high ($r_{jk0} >> 0$) but declined substantially due to competition from computer-based technology ($\Delta r_{jk} << 0$), we expect the slope coefficient b_j to be negative. Since b_j is in practice an elasticity—a one percent higher base-period wage will result in b_j-percent higher wage growth—it will capture changes in within-occupation wage differentials. Hence, TBTC implies a fall in within-occupation wage dispersion in routine-based occupations. But in occupations where routine skills are largely unimportant, we do not expect such technological changes to influence within-occupation wage differentials, since the oversupply of routine skills has already pushed its returns towards zero for all workers. By the same reasoning, we expect b_j to be positive for occupations where the price of "abstract skills" is high, if TBTC indeed increases the productivity of such skills. Since "service skills" should not be directly affected by TBTC, we would expect the slope coefficient to be close to zero or perhaps positive for typical service occupations.

Turning next to the intercept a_j, we can without loss of generality normalize the base period wage in each occupation to have a mean zero and by this write the intercept as

(6) $$a_j = \Delta\theta_j + \sum_{k=1}^{K} \Delta r_{jk} \bar{S}_{jk} \, .$$

The intercept, like the slope parameter b_j, depends on changes in the return to skill components, Δr_{jk}, but it also depends on $\Delta\theta_j$, which reflects changes in occupational wage differentials unrelated to skills. Note also that by normalizing the base period wage to have a mean zero, estimated intercepts will reflect average wage growth in each occupation, and different intercepts across occupations will therefore translate into changes in between-occupation wage differentials.

Our estimation proceeds in two steps. In the first we estimate equation (4) separately for each three-digit SSYK occupation by using individuals who stay in the same occupation over time (either 1990–2005 or 1975–1990). We use occupations rather than occupation-industry combinations (our definition of 'jobs' used before) to increase the sample sizes and thereby reduce the noise in the estimates of (4). In the second step, we test if the estimates are consistent with the predictions from TBTC by regressing the resulting slope coefficients and the intercepts on our occupation-specific measures of *abstract*, *routine*, and *service*. If TBTC is a main determinant of relative labor demand, we primarily expect to see positive estimates for *abstract* and negative estimates for *routine*. However, it should be noted that, according to our underlying wage model, estimates pertaining to the slope coefficients (within-occupation changes) are likely a better test of TBTC than corresponding estimates for intercepts (between-occupation changes) since a host of factors can potentially account for between-occupation changes, as reflected by $\Delta\theta_j$ in equation (6).

Some caveats should be recognized before proceeding to the results. The empirical approach relies on the assumption that the Roy-type linear skill pricing model of Firpo, Fortin, and Lemieux (2011) is a roughly correct description of wage setting in Sweden. Given the institutional setting in Sweden, with more or less coordinated wage settings between union and employer organizations across industries and occupations, this can obviously be questioned. Another drawback is the limited sample size that results from the additional restrictions on the sample. First, to be included, individuals must have employment in both 1990 and 2005 (or in 1975 and 1990). By construction, the sample then only consists of individuals who are less than 50 years old in the first year (since 64 years old is the maximum age in our analysis). Second, individuals have to be in the same three digit SSYK occupation in both years, that is, in the same occupation over a time span of 15 years. Third, to be able to get any meaningful estimates of equation (4), we demand that each occupation should have at least 10 longitudinal observations—that is, at least 10 individuals should be observed in both 1990 and 2005 (or in 1975 and 1990). As a result, we are left with approximately 27,500 individuals in our panel sample for each period. Since this is by no means a random sample, our estimates are potentially exposed to sample selection bias, and the direction of this potential bias is unclear. Because of the potential drawbacks with our panel approach, we for completeness also present results based on Firpo, Fortin, and Lemieux's (2011) cross-sectional implementation; this is further discussed below. A clear advantage with our panel approach, however, is that the results are not driven by individual movement across occupations or changes in the sample composition over time, as might be the case in a cross-sectional analysis.

Turning to the results from the panel approach, Table 6 contains the estimates from the second step, where we regress the estimated slope coefficients and the intercepts on our occupation-specific measures of *abstract*, *routine*, and *service*. For 1990–2005, there is mixed statistical support for TBTC. While the estimates for within-occupation wage differentials are in line with TBTC and strongly statistically significant, the null of a zero effect for between-occupation wage differentials cannot be rejected, with the exception of *abstrac*t in the first specification. The statistically significant point estimates for the slope coefficient do however imply a rather modest economic effect, in absolute terms, on within-occupation wage differentials. For instance, column (6) implies that in occupations with above average scores on *abstract*, a within-occupation wage differential of 10 percent in 1990 expanded by 0.89 percentage points up to 2005, to 10.89 percent. The corresponding number for routine is a decline of 0.54 percentage points. On the other hand, R^2 suggests that tasks indeed have been important for the *actual* changes in within-occupation wage differentials in Sweden between 1990 and 2005, at least as captured by the slope coefficient from the first stage. According to R^2, we are, with just three task dummies, able to explain 23 percent of the actual changes in slope coefficients; see column (4).

The last two columns of Table 6 contain results for 1975–1990 (additional results for this period are available on request). Overall, the estimates do not offer any statistical support for TBTC as a main determinant of wages in Sweden over this period. This is not surprising, since previous research (e.g., Hibbs, 1990; Edin and Topel, 1997) points toward the rise and fall of union-driven egalitarian wage policies, rather than market forces *per se*, as the main determinant of wage differentials in Sweden during the 1970s and 1980s.

Table 6: OLS on estimated intercepts (a_j) and slope coefficients (b_j) from equation (2)

	1990-2005						1975-1990	
	a_j			b_j			a_j	b_j
	(1)	(2)	(3)	(4)	(5)	(6)	(1)	(2)
Abstract	0.058**	0.035	0.036	0.069***	0.089***	0.089***	0.033	0.061
	(0.024)	(0.037)	(0.039)	(0.019)	(0.024)	(0.024)	(0.061)	(0.047)
Routine	-0.020	-0.012	-0.006	-0.045**	-0.052***	-0.054***	-0.066	0.017
	(0.025)	(0.031)	(0.033)	(0.017)	(0.019)	(0.020)	(0.052)	(0.040)
Service	0.002	-0.008	-0.003	-0.033**	-0.024	-0.026*	0.043	-0.063
	(0.017)	(0.022)	(0.021)	(0.014)	(0.014)	(0.014)	(0.050)	(0.039)
Education		0.009	0.008		-0.008	-0.007	-0.000	0.001
		(0.013)	(0.014)		(0.007)	(0.007)	(0.010)	(0.009)
Highly offshorable			0.045			-0.018	0.017	-0.066
			(0.040)			(0.027)	(0.062)	(0.058)
Obs.	94	94	94	94	94	94	83	83
R^2	0.153	0.169	0.183	0.230	0.240	0.242	0.134	0.165

Note: Dependent variables are occupation-specific estimates of the intercept and slope in equation (2). Standard errors are in parentheses.

* p<0.10, ** p<0.05, *** p<0.01

As mentioned above, we have also performed the same estimations as in Table 6 but instead used Firpo, Fortin, and Lemieux's (2011) cross-sectional implementation, where they, for each occupation, use cross-sectional wages at each decile in the estimation of (4) instead of individual wages; the reader is referred to the original study for a detailed account of this implementation. A clear advantage with this approach is obviously that we base our estimates on representative samples; the resulting samples contain over 120,000 individual for 1990–2005 and 1975–1990. We are also able to include 10 and 6 additional occupations for these two periods, respectively, as these now meet the requirement of having at least ten observations in both years. However, the cross-sectional approach forces us to impose distributional assumptions on the residual in equation (4) in order to get the desired interpretation of the estimated intercepts and slope coefficients; see Firpo, Fortin, and Lemieux (2011) for details. As in their paper, we therefore assume that changes in the distribution of the residual are similar across occupations, and that this can be fully accounted for by including quintile fixed effects (common for all occupations) in the first-stage estimations. Of course, we also have all the potential problems with job mobility and sample selection pertaining to cross-sectional data, as discussed above.

The results from the cross-sectional implementation are presented in Table 7. The conclusions from this analysis are—perhaps somewhat surprisingly—very similar to those from the panel analysis. There is statistical support for TBTC between 1990 and 2005 when it comes to within-occupation wage differentials but not for between-occupation wage differentials. Like in the panel analysis, R^2 imply that tasks have been important for changes in within-occupation wage differentials over this period. Also, results for 1975–1990 are again not supportive of TBTC. In our view, these similarities between the cross-sectional and panel implementation adds credibility to the obtained results.[29]

[29] This also indicates that, at least in our data, the repeated cross-section analysis used due to data limitations by Firpo, Fortin and Lemieux (2011) does not seem to cause any serious biases.

Table 7: OLS on estimated intercepts (a_j) and slope coefficients (b_j) from equation (2); cross-sectional implementation

	1990-2005						1975-1990	
	a_j			b_j			a_j	b_j
	(1)	(2)	(3)	(4)	(5)	(6)	(1)	(2)
Abstract	-0.017	0.046*	0.047*	0.158***	0.187***	0.187***	-0.044	0.047
	(0.024)	(0.026)	(0.027)	(0.042)	(0.075)	(0.074)	(0.045)	(0.041)
Routine	0.026	0.012	0.016	-0.104**	-0.111**	-0.120**	0.022	-0.012
	(0.024)	(0.021)	(0.022)	(0.040)	(0.045)	(0.049)	(0.056)	(0.060)
Service	-0.029	-0.004	-0.003	-0.031	-0.019	-0.024	0.069	-0.137**
	(0.028)	(0.025)	(0.025)	(0.035)	(0.038)	(0.036)	(0.058)	(0.058)
Education		-0.027***	-0.027***		-0.012	-0.011	-0.014	-0.004
		(0.008)	(0.008)		(0.026)	(0.025)	(0.012)	(0.011)
Highly offshorable			0.022			-0.064	0.059	-0.025
			(0.025)			(0.064)	(0.044)	(0.050)
Obs.	104	104	104	104	104	104	89	89
R^2	0.057	0.122	0.124	0.272	0.277	0.283	0.142	0.241

Note: Dependent variables are occupation-specific estimates of the intercept and slope in equation (2). Standard errors are in parentheses.

* p<0.10, ** p<0.05, *** p<0.01

IV. Concluding remarks

This paper investigates the connection between the Swedish wage profile of net job creation and the recently proposed hypothesis of 'task-biased technological change' (TBTC). We first document a pattern of job polarization between 1975 and 2005 with increased employment shares for the highest and lowest paid jobs. Unlike the polarization after 1990, the observed pattern of net job creation between 1975 and 1990 is however associated with a great deal of statistical uncertainty and is mostly accounted for by public sector employment.

In the next step of the analysis, we employ regressions to investigate the potential link between job tasks and the observed job polarization. Consistent with TBTC, we find that differences in the extent of routine versus non-routine tasks across jobs can potentially account for 44 percent of the observed job polarization between 1990 and 2005. The potential link between tasks and job creation prior to the 1990s is however much weaker.

We end the analysis with an investigation of wage changes and TBTC. Results for the period 1975–1990 are again unsupportive of TBTC. It is however harder to draw firm conclusions for the period 1990–2005. While changes in between-occupation wage differentials do not lend statistical support to TBTC, changes in within-occupation wage differentials fit nicely with predictions from TBTC.

Overall, while our results for the period 1975–1990 do not contain any clear evidence in favor of TBTC, results for the more recent period 1990–2005 are more supportive, although not conclusive. One the one hand, the notable job polarization during this period and its strong link to routine versus non-routine tasks across jobs add support to the notion of TBTC. On the other hand, our study of between-occupation wage differentials casts some doubt on the impact of TBTC in Sweden during this period, and this doubt is certainly amplified by similar findings for Germany during this period in Kampelmann and Rycx (2011). As such, our results for Sweden indicate that it might be premature to treat TBTC as a stylized fact across all OECD countries; see also Fernández-Macías and Hurley (2008) for cross-country evidence in this direction.

An explanation for the lack of a clear link between between-occupation wage differentials and predictions from TBTC in Sweden and Germany could of course be the more regulated and coordinated wage setting in these countries compared to countries such as the U.S. One could perhaps imagine that Swedish wage rigidities cause TBTC to primarily affect employment and unemployment risks across different jobs and workers, rather than their wages.[30] An important topic for future research is thus to investigate changes in unemployment and unemployment risks in connection to the hypothesis of TBTC.

References

Acemoglu, D. (2001), God Jobs versus Bad Jobs, *Journal of Labor Economics* 19, 1–21.

Acemoglu, D. (2003), Cross-Country Inequality Trends, *Economic Journal* 113, 121–149.

Acemoglu, D. and Autor, D. (2011), Skills, Tasks and Technologies: Implications for Employment and Earnings, in O. Ashenfelter and D. Card (eds.), *Handbook of Labor Economics Volume 4*, Elsevier Science B.V., Amsterdam.

Angrist, J. and Pischke, J.-S. (2009), *Mostly Harmless Econometrics*, Princeton University Press, Princeton.

[30] Since the Swedish wage bargaining has moved towards the firm- or individual level over the last two decades, especially for white collar workers (e.g., Lundborg, 2005), the extent of Swedish wage rigidities is an open question.

Atkinson, A. (2008), *The Changing Distribution of Earnings in OECD Countries*, Oxford University Press, Oxford.

Autor, D. and Dorn, D. (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., Katz, L. and Kearney, M. (2008), Trends in U.S. Wage Inequality: Re-Assessing the Revisionists, *Review of Economics and Statistics* 90, 300–323.

Autor, D., Levy, F. and Murnane, R. (2003), The Skill Content of Recent Technological Change: An Empirical Exploration, *Quarterly Journal of Economics* 118, 1279–1333.

Bickel, P., Götze, F. and van Zwet, W. (1997), Resampling Fewer than n observations: Gains, Losses and Remedies for Losses, *Statistica Sinica* 7, 1–31.

Bickel, P. and Sakov, A. (2008), On the Choice of m in the m Out of n Bootstrap and Confidence Bounds for Extrema, *Statistica Sinica* 18, 967–985.

Björklund, A. and Freeman, R. (2010), Searching for Optimal Inequality/Incentives, in R. Fremman, B. Swedenborg and R. Topel (eds.), *Reforming the Welfare State: Recovery and Beyond in Sweden*, The University of Chicago Press, Chicago.

Blinder, A. (2009), How Many U.S. Jobs Might be Offshorable?, *World Economics* 10, 41–78.

Blinder, A. and Krueger, A. (2013), Alternative Measures of Offshorability: A Survey Approach, *Journal of Labor Economics* 31, S97–S128.

Cahuc, P. and Zylberberg, A. (2004), *Labor Economics*, MIT Press, Cambridge.

Cameron, C. and Trivedi, P. (2005), *Microeconometrics: Methods and Applications*, Cambridge University Press, New York.

Domeij, D. (2008), Rising Earnings Inequality in Sweden: The Role of Composition and Prices, *Scandinavian Journal of Economics*, 110, 609–634.

Dustmann, C. and Meghir, C. (2005), Wages, Experience and Seniority, *Review of Economic Studies* 72, 77–108

Dustmann, C., Ludsteck, J. and Schönberg, U. (2009), Revisiting the German Wage Structure, *Quarterly Journal of Economics* 124, 843–881.

Edin, P.-A. and Fredriksson, P. (2000), LINDA – Longitudinal Individual Data for Sweden, Working Paper 2000:19, Department of Economics, Uppsala University.

Edin, P.-A. and Holmlund, B. (1995), The Swedish Wage Structure: The Rise and Fall of Solidarity Wage Policy?, in R Freeman and L Katz (eds.), *Differences and Changes in Wage Structures*, University of Chicago Press, Chicago.

Edin, P.-A. and Topel, R. (1997), Wage Policy and Restructuring: The Swedish Labor Market Since 1960, in R. Freeman, R. Topel and B. Swedenborg (eds.), *The Welfare State in Transition*, The University of Chicago Press, Chicago.

Fernández-Macías, E. and Hurley, J. (2008), More and Better Jobs: Patterns of Employment Expansion in Europe, The European Foundation for the Improvement of Living and Working Conditions, ERM Report 2008.

Firpo, F., Fortin, N. and Lemieux, T. (2011), Occupational Tasks and Changes in the Wage Structure, Discussion Paper 5542, IZA.

Goos, M. and Manning, A. (2007), Lousy and Lovely Jobs: The Rising Polarization of Work in Britain, *Review of Economics and Statistics* 89, 119–133.

Goos, M., Manning, A. and Salomons, A. (2009), Job Polarization in Europe, *American Economic Review: Papers & Proceedings* 9, 58–63.

Goos, M., Manning, A. and Salomons, A. (2010), Recent Changes in the European Employment Structure: The Roles of Technology and Globalization, mimeo, University of Leuven.

Gustavsson, M. (2006), The Evolution of the Swedish Wage Structure: New Evidence for 1992–2001, *Applied Economics Letters*, 13, 279–286.

Hibbs, D. (1990), Wage Dispersion and Trade Union Action in Sweden, in I. Persson (ed.), *Generating Equality in the Welfare State—The Swedish Experience*, Norwegian University Press, Oslo.

Holmlund, B. (2006), The Rise and Fall of Swedish Unemployment, in M. Werding (ed.), *Structural Unemployment in Western Europe: Reasons and Remedies*, MIT Press, Cambridge.

Horowitz, J. (2001), The Bootstrap, in J. Heckman and E. Leamer (eds.), *Handbook of Econometrics Volume 5*, Elsevier Science B.V., Amsterdam.

Katz, L. and Murphy, K. (1992), Changes in Relative Wages, 1963–87: Supply and Demand Factors, *Quarterly Journal of Economics* 107, 35–78.

Katz, L. and Autor, D. (1999), Changes in the Wage Structure and Earnings Inequality in O. Ashenfelter and D. Card (eds), *Handbook of Labor Economics*, Vol. 3A, Elsevier Science B.V., Amsterdam.

Kampelmann, S. and Rycx, F. (2011), Task-Biased Changes of Employment and Remuneration: The Case of Occupations, Discussion Paper 5470, IZA.

Kennedy, P. (1998), *A Guide to Econometrics*, Blackwell, Oxford.

Rosen, S. (1997), Public Employment, Taxes, and the Welfare State in Sweden, in R. Freeman, R. Topel and B. Swedenborg (eds.), *The Welfare State in Transition*, The

University of Chicago Press, Chicago.

Statistics Sweden (1998), *Standard Classification of Occupations 1996. Reports on Statistical Co-ordination for the Official Statistics of Sweden 1998:3*, Stockholm.

Wooldridge, J. (2010), *Econometric Analysis of Cross Section and Panel Data*, The MIT Press, Cambridge.

Wright, E. O. and Dwyer, R. (2003), The Patterns of Job Expansion in the USA: A Comparison of the 1960s and 1990s, *Socio-Economic Review* 1, 289–325.

Åberg, R. (2004), Vilka jobb har skapats på den svenska arbetsmarknaden de senaste decennierna?, *Ekonomisk debatt* 32(7), 37–46.

Appendix

Table A1: Summary statistics

	All			Private sector			Public sector		
	Mean	St. dev.	Obs.	Mean	St. dev.	Obs.	Mean	St. dev.	Obs.
Share female									
1975	0.38	0.37	109,558	0.30	0.32	76,803	0.59	0.39	32,755
1990	0.40	0.36	108,030	0.31	0.31	76,133	0.63	0.36	31,897
2005	0.39	0.34	104,125	0.28	0.28	75,398	0.65	0.33	28,727
Years of schooling									
1975	9.9	2.0	109,544	9.3	1.4	76,794	11.2	2.5	32,750
1990	10.6	1.8	108,008	10.2	1.3	76,113	11.7	2.3	31,895
2005	11.6	1.4	104,125	11.3	1.0	75,398	12.4	1.9	28,727
Age									
1975	39.6	4.0	109,558	39.5	3.8	76,803	39.7	4.4	32,755
1990	39.8	3.6	108,030	39.2	3.5	76,133	41.2	3.4	31,897
2005	42.3	4.2	104,125	41.4	3.8	75,398	44.8	4.1	28,727
Share foreign-born									
1975	0.09	0.08	109,558	0.10	0.08	76,803	0.06	0.05	32,755
1990	0.09	0.08	108,030	0.09	0.08	76,133	0.08	0.06	31,897
2005	0.10	0.08	104,125	0.11	0.09	75,398	0.09	0.07	28,727
Task importance, dummy									
Abstract	0.34	0.47	107,161	0.31	0.46	75,135	0.40	0.49	32,026
Routine	0.52	0.50	107,161	0.63	0.48	75,135	0.25	0.43	32,026
Service	0.52	0.50	107,161	0.43	0.49	75,135	0.74	0.43	32,026

Table A2: Correlations for task dummies and offshore dummies, job level

	Abstract	Routine	Service	Offshorable	Highly offshorable
Abstract	1				
Routine	-0.0967^{*}	1			
Service	0.0756	-0.678^{***}	1		
Offshorable	0.111^{*}	-0.0218	-0.0393	1	
Highly offshorable	0.216^{***}	-0.255^{***}	0.102^{*}	0.322^{***}	1

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table A3: Largest jobs in the private and public sectors in 1975

Employment	Job (*Occupation in Industry*)
	Private sector
163,564	Shop and stall salespersons and demonstrators in Wholesale and retail trade; repair of motor vehicles,
132,657	Building frame and related trades workers in Construction
122,359	Crop and animal producers in Agriculture, hunting and forestry
91,736	Motor-vehicle drivers in Transport, storage and communication
73,930	Finance and sales associate professionals in Wholesale and retail trade; repair of motor vehicles, motorcycles
69,280	Managers of small enterprises in Wholesale and retail trade; repair of motor vehicles, motorcycles and
56,859	Agricultural and other mobile-plant operators in Agriculture, hunting and forestry
56,350	Building finishers and related trades workers in Construction
51,280	Metal- and mineral-products machine operators in Manufacture of machinery and equipment n.e.c.
49,271	Other office clerks in Wholesale and retail trade; repair of motor vehicles, motorcycles and personal and
	Public sector
667,920	Personal care and related workers in Public administration and defense; compulsory social security
216,261	Primary education teaching professionals in Real estate, renting and business activities
122,585	Other office clerks in Financial intermediation
109,188	Nursing associate professionals in Public administration and defense; compulsory social security
76,943	Public service administrative professionals in Financial intermediation
72,444	Armed forces in Financial intermediation
64,667	Other office clerks in Public administration and defense; compulsory social security
60,153	Mail carriers and sorting clerks in Wholesale and retail trade; repair of motor vehicles, motorcycles and
55,579	Stores and transport clerks in Wholesale and retail trade; repair of motor vehicles, motorcycles and personal
55,049	Electrical and electronic equipment mechanics and fitters in Wholesale and retail trade; repair of motor

Note*:* Employment numbers are rounded full-time equivalents, and are rescaled to match aggregate

employment for the whole of Sweden.

Table A4: Most growing and most shrinking jobs in each quintile, 1975-2005

Δ Employment	Job (*Occupation* in *Industry*)
Quintile 1	
145,487	Personal care and related workers in Health and social work
86,768	Personal care and related workers in Education
21,769	Pre-primary education teaching associate professionals in Education
-30,513	Textile-, fur- and leather-products machine operators in Manufacture of textiles and textile products
-32,942	Agricultural and other mobile-plant operators in Agriculture, hunting and forestry
-82,742	Crop and animal producers in Agriculture, hunting and forestry
Quintile 2	
43,553	Pre-primary education teaching associate professionals in Health and social work
30,479	Shop and stall salespersons and demonstrators in Wholesale and retail trade; repair of motor vehicles, motorcycles and personal and household goods
20,258	Office secretaries and data entry operators in Health and social work
-18,554	Other office clerks in Wholesale and retail trade; repair of motor vehicles, motorcycles and personal and household goods
-31,762	Forestry and related workers in Agriculture, hunting and forestry
-32,878	Other office clerks in Public administration and defence; compulsory social security
Quintile 3	
39,743	Secondary education teaching professionals in Health and social work
25,050	Health associate professionals (except nursing) in Health and social work
12,231	Protective services workers in Real estate, renting and business activities
-22,221	Metal- and mineral-products machine operators in Manufacture of machinery and equipment n.e.c
-28,863	Managers of small enterprises in Wholesale and retail trade; repair of motor vehicles, motorcycles and personal and household goods
-33,554	Building frame and related trades workers in Construction
Quintile 4	
38,129	Nursing and midwifery professionals in Health and social work
30,336	Nursing associate professionals in Health and social work
25,041	Psychologists, social work and related professionals in Health and social work
-13,484	Craft printing and related trades workers in Manufacture of pulp, paper and paper products; publishing and printing
-14,596	Stores and transport clerks in Transport, storage and communication
-14,775	Metal moulders, welders, sheet-metal workers, structural-metal preparers and related trades workers in Manufacture of machinery and equipment n.e.c.
Quintile 5	
66,340	Primary education teaching professionals in Health and social work
58,277	Computing professionals in Real estate, renting and business activities
27,782	Health professionals (except nursing) in Health and social work
-14,374	Physical and engineering science technicians in Manufacture of machinery and equipment n.e.c.
-19,237	Physical and engineering science technicians in Construction
-33,614	Primary education teaching professionals in Education

Note*:* Quintile 1 is the lowest wage quintile Changes in employment is rounded full-time equivalent employment, and are rescaled to match aggregate changes for the whole of Sweden.