

Graetz, Georg

Working Paper

Technological change and the Swedish labor market

Working Paper, No. 2020:19

Provided in Cooperation with:

IFAU - Institute for Evaluation of Labour Market and Education Policy, Uppsala

Suggested Citation: Graetz, Georg (2020) : Technological change and the Swedish labor market, Working Paper, No. 2020:19, Institute for Evaluation of Labour Market and Education Policy (IFAU), Uppsala

This Version is available at:

<https://hdl.handle.net/10419/246028>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Technological change and the Swedish labor market

Georg Graetz

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

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

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

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

ISSN 1651-1166

Technological change and the Swedish labor market^a

Georg Graetz^b

September 30, 2020

Abstract

This report takes stock of recent research into the effects of technology on the labor market; assesses to what extent the Swedish labor market has been affected by technological change in the past three decades, in particular with respect to the themes highlighted by the research; and draws lessons for the future. The main conclusion is that so far, the Swedish labor market has adapted well to technological change, as it exhibits stable employment rates and steady wage growth across the distribution. Furthermore, there are no clear signs that this situation should change in the coming decades.

Key words: Technology, employment, wages, the Swedish economy

JEL codes: J21, J23, J24, J31, O30, O52

^aThis is a slightly revised version of a report prepared for the Swedish Fiscal Policy Council (Graetz, 2020). I thank the members and staff of the Council as well as Adrian Adermon and Anders Åkerman for helpful comments and suggestions, Axel Malmcrona for excellent research assistance, and Lena Hensvik and Oskar Nordström Skans for sharing data. Any remaining errors are my own.

^bDepartment of Economics, Uppsala University, and IFAU. Email: georg.graetz@nek.uu.se.

Contents

1	Introduction	3
2	From skills to tasks: a review of the literature	4
2.1	The rise in the college premium in the US and the case for skill-biased technological change	5
2.2	Job polarization and the task-level impact of new technology	7
2.3	The fall in the labor share	10
2.4	Evidence on automation: industrial robots	13
2.5	Summary	15
3	Swedish labor market trends 1985-2017	16
3.1	Trends in employment and composition	16
3.2	Trends in wage inequality and skill returns	20
3.3	Summary and interpretation of employment and inequality trends	25
4	The occupation-level and regional impacts of technological change	26
4.1	The career costs of occupational decline	26
4.2	Task-biased technological change and occupational wage inequality	28
4.3	The impact of technological change across regions	31
5	The future	34
6	Conclusions	39
	Appendix figures and tables	53

1 Introduction

Fears that new technologies are about to put many if not most people out of work are once again a topic of conversation among academics, policy makers, and the public (Autor, 2015; Shiller, 2019). Economists agree that since the beginning of the Industrial Revolution, technology has largely been a blessing, causing a spectacular rise in incomes and standards of living (Jones, 2016). And while there were periods when many were adversely affected by technology-induced disruption without gaining improved job opportunities or higher wages—especially in the first half of the 19th century (Frey, 2019)—in the medium and long term technological change has indeed brought about sustained growth in wages across the distribution, as well as job growth that has more than kept up with increases in population. However, recent advances in robotics and artificial intelligence lead some to claim that this time is different—prospects for less-skilled workers may be deteriorating, as automation threatens to proceed at a higher pace (Brynjolfsson and McAfee, 2014; Ford, 2015), and the creation of new tasks appears to slow down (Acemoglu and Restrepo, 2019b). Others suggest that the demand for middle-skill workers—who have lost out from recent technological change—may well pick up again (Autor, 2015). Yet others question the ability of modern technologies such as machine learning and robotics to deliver sustained productivity growth (Gordon, 2012, 2014).

My aim in this report is to take stock of recent research into the effects of technology on the labor market; to assess to what extent the Swedish labor market has been affected by technological change in the past three decades, in particular with respect to the themes highlighted by the research; and to draw lessons for the future.

The Swedish economy has seen steady growth after recovering from the deep recession in the early 1990s, and has adopted new technologies such as information and communication technologies (ICT) and industrial robots at rates comparable to other developed countries (O’Mahony and Timmer, 2009; Graetz and Michaels, 2018). It is also known to have a vibrant start-up culture (Semuels, 2017). Finally, it is no exception in terms of sectoral shifts and occupational polarization. Thus, Sweden appears to be part of the group of economies that continue to push the technological frontier, and the question of technology’s impact on the labor market

is no less pressing here than elsewhere.¹

I begin the report by reviewing recent research into the impact of technological change on the labor market in Section 2, with a focus on the experience of the US and other major economies, as well as on the evolution of theoretical frameworks to account for a variety of salient phenomena.

I then present a series of Swedish labor market trends in Section 3, including the employment rate, the sectoral and occupational distribution of employment, and various measures of wage inequality, covering the period 1985-2017. I connect these findings to the experience of other countries as well as the theoretical frameworks discussed in Section 2.

In Section 4, I report on ongoing research that explores the occupation-level impact of technology in Sweden, including the consequences of occupational decline for workers' careers, and the impact of technological change on the occupational wage structure. I also explore how technology adoption varies across Swedish municipalities, and how this regional variation is associated with differences in employment growth, skill mix changes, and changes in industrial composition.

In Section 5, I review recent research into the future impact of technology, and how it applies to Sweden. Finally, I summarize my results and draw conclusions in Section 6.

2 From skills to tasks: a review of the literature

The impact of technological change on the economy is complex, difficult to predict, and often hard to grasp even in hindsight. There are perhaps only two statements that can be made with confidence: first, that new technologies increase productivity, and second, that they cause disruption. Beyond that, there is much uncertainty. Technological change may favor the skilled over the un-skilled or vice versa, so that inequality may increase or decline (perhaps even evolve

¹Sweden is also no exception in terms of openness to trade, and has experienced increased import competition from China and other low-wage countries since the mid-1990s. A discussion of the effects this had on the Swedish labor market is beyond the scope of this report. There is evidence that import competition raised high-skill wages while having no effect on the wages of the less-skilled (Baziki, 2015); and that ICT facilitates the reallocation of workers across firms in response to increased import competition (Baziki, Ginja, and Milicevic, 2018). Further important demographic and labor market trends in Sweden include a high rate of immigration as well as continued urbanization. The interactions of these trends with technological change is an important but rather under-studied area of research (Autor, 2019), and a discussion is beyond the scope of this paper.

in different directions in different parts of the distribution); automation may proceed unevenly across occupations; new jobs are created even as technology makes many jobs obsolete, with the net effect being unclear; the adoption of new technologies may require organizational changes, or a geographic reallocation of economic activity; finally, technology can affect market size, which in turn affects inequality.

A review of the research into technological change over the past 30 years confirms this wide range of possibilities. The following sub-sections trace out in roughly chronological order how the emphasis of the research has shifted, always motivated by new emerging phenomena. Section 2.1 focuses on the relationship between technological change and the skill premium. Section 2.2 discusses the uneven impact of new technologies across job tasks, and the resulting shifts in occupational employment. Section 2.3 reviews recent empirical work on the labor share, and asks to what extent its evolution may be driven by technology. And Section 2.4 investigates whether automation implies job losses, both empirically and theoretically.

2.1 The rise in the college premium in the US and the case for skill-biased technological change

Between 1980 and 2008, the college premium in the US doubled (Acemoglu and Autor, 2011). While in 1980, a worker in possession of a four-year college degree earned on average about 50 percent more than a comparable worker (in terms of gender, age, and race) with no more than a 12-year high school education, in 2008 the figure stood at nearly 100 percent. At the same time, college-educated labor had become relatively more abundant. In a basic supply-and-demand framework, this increased abundance by itself would imply a *decrease* in the college premium. The fact that the premium instead increased suggests that demand for college-educated workers has shifted up.^{2,3}

There is a near-consensus among economists that the computer revolution—the spread of

²It is not necessary to assume a frictionless, perfectly competitive labor market, where each type of labor is paid its marginal product, to obtain this prediction. What is required is that in the medium to long run, wages respond to shifts in supply and demand in qualitatively the same way as in the frictionless, competitive benchmark. Many models of the labor market, including those assuming search frictions, collective bargaining, or some other form of market power, will satisfy this requirement.

³There is some evidence that since 2008, the demand for skilled labor has not increased further, or even declined somewhat (Beaudry, Green, and Sand, 2016; Autor, Goldin, and Katz, 2020).

information and communication technologies (ICT)—is the main factor behind the increased demand for skills.⁴ The idea is that ICT benefits skilled workers disproportionately, by making them more productive in tasks such as data processing, graphic design, or monitoring and managing production processes. Models of such skill-biased technological change (SBTC) perform well in explaining US data on relative quantities and wages of college-educated labor (Katz and Murphy, 1992; Krusell, Ohanian, Ros-Rull, and Violante, 2000).

Beyond time-series evidence, many aspects of cross-industry and individual-level data from the US are consistent with SBTC. Industries that adopted computers the fastest also saw the largest increase in the wage bill shares of skilled workers (Autor, Katz, and Krueger, 1998). College-educated workers use computers at a higher rate than high-school graduates, and workers who use computers earn higher wages than comparable workers who do not (Krueger, 1993). While one may worry that other factors could give rise to the same patterns,⁵ SBTC has received further support from studies that take advantage of natural experiments to estimate the causal effects of ICT on the skill premium. Akerman, Gaarder, and Mogstad (2015) study the quasi-exogenous roll-out of broadband internet in Norway and find that internet access causes an increase in the skill premium via a greater skill intensity of firms' production processes. Gaggl and Wright (2017) leverage a tax credit in the UK for causal identification and similarly find that ICT adoption benefits skilled labor.⁶

Most developed countries have seen trends in skill premia and the supply of skilled labor similar to the ones in the US discussed here, and industry-level skill-upgrading follows highly similar patterns across countries (Berman, Bound, and Machin, 1998), well in line with SBTC (Machin and Van Reenen, 1998). I discuss the case of Sweden in Section 3 below.

The theory of SBTC nevertheless has two important weaknesses. First, the theory is not explicit about how exactly ICT complements skilled labor, in particular with respect to the task content of skilled work, and it is silent on how the impact of ICT varies across occupations.

⁴See Acemoglu (2002), Autor, Katz, and Kearney (2008), Acemoglu and Autor (2011) in favor of this view. DiNardo, Fortin, and Lemieux (1996) and Card and DiNardo (2002) put forward alternative views.

⁵For instance, DiNardo and Pischke (1997) document that there is also a wage premium associated with the use of pencils.

⁶Both Akerman, Gaarder, and Mogstad (2015) and Gaggl and Wright (2017) place their findings in the context of not only the SBTC literature, but also the one on task-biased technological change to be discussed in Section 2.2, providing support for key empirical predictions put forward in both literatures.

Second, SBTC predicts that low-skilled labor loses out in relative terms, but cannot rationalize the fall in the *level* of these workers' real wages, as has occurred in the US. It turns out that these weaknesses are related and have a common cure, as is explained in the next sub-section.

2.2 Job polarization and the task-level impact of new technology

Turning one's attention to the evolution of employment across occupations, and given the empirical success of SBTC just discussed, one may expect a positive relationship between occupational employment growth and the skill intensity of occupations. However, the typical pattern is instead a more nuanced one whereby not only high-skill but also low-skill occupations (measured either in terms of wages or educational attainment) have gained employment shares at the expense of the middle. This phenomenon, called job polarization, was first documented for the US (Wright and Dwyer, 2003; Autor, Katz, and Kearney, 2006) and the UK (Goos and Manning, 2007), and subsequently confirmed to hold for the vast majority of European economies (Goos, Manning, and Salomons, 2014), including Sweden (Adermon and Gustavsson, 2015).

Basic models of skill-biased technological change cannot explain job polarization, which is not surprising as they were not designed for this purpose—indeed, these models do not contain occupations and hence are silent on occupational employment shifts by construction. Job polarization can however be explained as resulting from technological change within the *task framework* developed by Autor, Levy, and Murnane (2003), Acemoglu and Autor (2011), and Acemoglu and Restrepo (2018a). In this framework, production requires the completion of a large number of distinct tasks. Individual workers differ in their ability to perform each task, and in equilibrium specialize in the task (or set of tasks) in which they have a comparative advantage. The framework can readily incorporate automation by allowing machines to be capable of performing a subset of tasks, at a cost sufficiently low so firms are willing to use them. Task-biased technological change (TBTC) refers to increased automation, that is, an expansion of the set of tasks performed by machines.⁷

For TBTC to generate job polarization, it must be that machines' comparative advantage

⁷This may happen because machines become capable of performing a wider set of tasks. Alternatively, it may be the result of a falling rental price of machines inducing firms to automate tasks which they could have automated even before, but did not do so because it would not have been profitable.

is in tasks originally performed by middle-skill workers. This is in fact a reasonable scenario given the experience of researchers working on artificial intelligence (AI). As Moravec (1988) observes, “it is comparatively easy to make computers exhibit adult level performance on intelligence tests or playing checkers, and difficult or impossible to give them the skills of a one-year-old when it comes to perception and mobility.” Polanyi (1966) stresses the importance of tacit knowledge: “We can know more than we can tell.” We may be capable of having a conversation or riding a bicycle, but at the same time are unable to break down these activities into simple, codifiable rules. The latter however is a critical requirement for automation (at least until the recent advent of machine learning), as highlighted by Autor, Levy, and Murnane (2003). In their terminology, tasks that could be automated following the arrival of ICT are called *routine*—repetitive, predictable, and codifiable tasks such as record keeping, numerical calculations, or the assembly of manufacturing goods. In contrast, non-automatable *non-routine* tasks included both low-skill tasks such as waiting tables and driving, as well as high-skill tasks such as managing, making decisions based on a variety of data, and arguing a legal case. Thus, job polarization results from TBTC because routine tasks, including both clerical office tasks as well as manufacturing work, tended to be performed by middle-skill workers.

To take the task framework to the data, researchers have sought to characterize the task content of occupations. The seminal contribution is Autor, Levy, and Murnane (2003), who classify occupations according to their content of routine and non-routine tasks (among other dimensions), and document a positive correlation between computer adoption and routine task content, as well as a decline in employment in routine-intensive occupations since 1960. A more recent line of research characterizes changes in routine task content within occupations over time, for instance using job descriptions from newspaper job ads, as in Atalay, Phongthiengtham, Sotelo, and Tannenbaum (forthcoming). They find that since 1950, not only has employment in the US shifted away from initially routine-intensive jobs, but also that routine tasks have become less common *within* occupations. This result highlights the distinction between tasks and occupations—in reality, occupations are bundles of tasks, and these bundles often change over time.

Furthermore, the task framework allows not only for automation of tasks and hence the replacement of labor, but also for the creation of new tasks and thus a “re-instatement effect” (Acemoglu and Restrepo, 2019a). Creation of new tasks may be measured by the incidence of new occupational titles when occupational classifications change (Acemoglu and Restrepo, 2018b), and the net effect of technological change—the difference between replacement and re-instatement effects—can be quantified under additional assumptions, as I discuss in the next sub-section.

The task framework has rich implications for the compensation of production factors and hence the income distribution. First, consider the distribution of wages (labor income). Because different skills may be differently productive across tasks, TBTC likely implies changes in the returns to skills. For instance, physical strength has become less valuable due to the mechanization of manufacturing. Analytical ability, in contrast, becomes more valuable as it is used in tasks that are not yet automated—such as decision-making in management—but that benefit from cheaper data processing and numerical calculation. In general, when tasks are complementary in production, automation benefits the workers performing the non-automated tasks. Indeed, the returns to cognitive skills have increased over the past half century, but interestingly, in recent decades it is social-interactive abilities that have become relatively more valuable (Deming, 2017; Edin, Fredriksson, Nybom, and Öckert, 2017).

Second, consider the *level* of the wage. In traditional producer theory it is usually not possible to generate a negative effect of technology on wages (at least under reasonable assumptions, see Caselli and Manning, forthcoming). In the task framework, the wage change resulting from automation consists of two components, a negative replacement effect and a positive productivity effect. The former is because workers compete for a smaller set of tasks. The latter is due to complementarity across tasks—when cheaper capital is producing complementary inputs, labor productivity in non-automated tasks increases. Thus, automation can lead to lower wages if the replacement effect dominates, which happens in the case of “so-so innovations” that induce firms to automate but do not deliver large productivity gains (Acemoglu and Restrepo, 2018a).

Finally, the task framework sheds light on the distribution of income across capital and

labor. I discuss this issue in detail in the following sub-section.

2.3 The fall in the labor share

In macroeconomics textbooks, the share of GDP accruing to labor is usually introduced as a quantity that has been stable over many decades or even centuries. However, recent research highlights a *downward* trend in the labor share across countries and industries since around 1980—about a 5-percentage-point decline for the global average (Karabarbounis and Neiman, 2014). The major economies that have seen declining labor share include for instance the US, Germany, and Japan, though in Sweden the share appears rather stable (Konjunkturinstitutet, 2018).

Trends in the labor share are generally somewhat difficult to pin down, for two reasons. First, the labor share tends to be quite volatile, displaying large short-to-medium-run fluctuations. This means that statements about long-run trends can be sensitive to the choice of start year. For instance, since the year 2000, the number of countries that saw declining labor shares was about equal to the number of countries with rising shares (Aum, Koh, and Santaaulalia-Llopis, 2019)—Sweden is among the latter group. Second, the definition of the labor share used by statistical agencies may vary across countries, and may even change within countries over time, especially with respect to the treatment of self-employment (Elsby, Hobijn, and Sahin, 2013; Hagelund, Nordbø, and Sauvik, 2017; Gutierrez and Piton, 2019).

Having noted these caveats, I accept for now the view that labor share changes since 1980 are best characterized as following a downward trend, and discuss potential explanations for this. One set of explanations relates to technology. Karabarbounis and Neiman (2014) argue that the strong decline in the price of equipment capital—due to technological advances such as personal computers (Nordhaus, 2007) and industrial robots (Graetz and Michaels, 2018)—has induced firms to substitute capital for labor. In their theoretical model, this results in a lower labor share if capital and labor are sufficiently substitutable in production—in technical terms, if the elasticity of substitution between capital and labor exceeds one. Karabarbounis and Neiman (2014) argue that this is a reasonable interpretation of the data, but acknowledge

that prior literature in most cases estimated this elasticity to be below one.

Alternatively, one can invoke the task framework to explain the labor share decline as resulting from technological change. The lower price of equipment capital, and enhanced capabilities of machinery, lead firms to use machines in a larger share of tasks. In task models, a production factor's task share is closely linked to the factor's share in total income, and this property does not hinge on the value of the substitution elasticity.⁸ Acemoglu and Restrepo (2019a) estimate that increased automation and a slower rate of new-task creation accounts for all of the decline in the labor share in the US. In Graetz (2019), I apply their methodology to European data and confirm this result, while also documenting a positive correlation in industry-level task content changes across countries. While the method depends on many strong assumptions, this evidence can be taken as suggestive of a link between technological change and labor share trends.

A second set of explanations for the falling labor share invokes firms' increased market power. In the case of perfect competition and constant-returns-to-scale production technologies,⁹ all income accrues to the factors of production, so that a falling labor share is equivalent to a rising capital share. However, if firms have market power (in the product or factor markets, or both) they will capture part of their value added in the form of pure economic profits. Profits can be difficult to distinguish from capital income in the data, however, because firms own much of the capital stock rather than renting it. Measuring capital income thus requires an estimate of the rental rate of capital, in addition to an estimate of the amount of capital used in production. Barkai (2017) conducts this exercise for the US and finds that the capital share has also declined, implying a rising share of pure profits. However, Karabarbounis and Neiman (2018) raise the possibility that rental prices have been mismeasured, and that the rise in the profit share is thus overstated.

An alternative strategy to track firms' market power is to estimate the markup they charge

⁸In the task framework, the substitution elasticity between capital and labor in the aggregate production function derives from the substitutability of different tasks, and might therefore be rather low. But even an elasticity well below one would imply that increased automation corresponds to a lower labor share.

⁹Under constant returns to scale, a doubling of all inputs leads to a doubling of output. With increasing returns, doubling inputs implies that output more than doubles, as can be the case with non-rival goods such as software, or when there are fixed costs to enter production. Increasing returns cannot support a perfectly competitive market, as firms would be making losses. With increasing returns, some form of market power—either in product or factor markets—is thus required.

over marginal cost. De Loecker, Eeckhout, and Unger (forthcoming) apply this strategy to firm-level data and find that markups in the US increased from 20 percent over marginal cost in 1980 to 60 percent today. De Loecker and Eeckhout (2018) document similar trends for most regions of the world. De Loecker, Eeckhout, and Unger (forthcoming) explore the implications of rising markups for the labor share, as do Eggertsson, Robbins, and Wold (2018). Basu (2019) argues that the estimated rise in markups implies a decline in the labor share that is much larger than what is found in the data. Traina (2018) and Karabarbounis and Neiman (2018) highlight measurement issues related to markup estimation using firm-level data, and argue that markups appear to be stable under alternative, not necessarily less defensible assumptions.

While there appears to be no consensus yet about recent trend in markups, researchers largely agree that there has been a rise in concentration—both in terms of employment and sales—in most industries at the national level in the US since the early 1980s (Rossi-Hansberg, Sarte, and Trachter, 2018). Increased concentration is driven by the largest firms in an industry becoming more dominant. Since these large firms tend to have lower payroll-to-sales ratios, increased concentration accounts for part of the fall in the labor share (Autor, Dorn, Katz, Patterson, and Van Reenen, forthcoming).¹⁰ Theoretically, it is not clear whether increased concentration should be interpreted as greater market power. On one hand, an increased dominance of large firms may be seen as synonymous with diminished competition. On the other hand, increased competitive pressures could cause smaller, less productive firms to go out of business, leaving only the large ones to survive. Such increased “toughness” (Autor, Dorn, Katz, Patterson, and Van Reenen, forthcoming) could be the result of technological changes allowing for greater scale economies, or increased international competition.

To summarize, there is no consensus yet whether labor shares have undergone pervasive decline, and if so, whether this would be due to technological change or due to increased market power. There is evidence that increased concentration may be driving a decline, but it is not

¹⁰A related issue which has recently received attention is local concentration of employment and the wage setting power of firms. While higher local concentration is associated with lower wages (Azar, Marinescu, Steinbaum, and Taska, 2018), local concentration has actually decreased in the US (Lipsius, 2018; Rinz, 2018). This appears to be driven by the same forces as the increase in concentration at the national level, as entry of large firms in a local labor market in fact leads to lower concentration (Rossi-Hansberg, Sarte, and Trachter, 2018) locally. Trends in *local* concentration and increased monopoly power thus cannot explain the falling labor share or rising inequality.

clear whether such increased concentration can be interpreted as resulting from less competitive pressures, so that a strong case for more aggressive antitrust action cannot be made.

2.4 Evidence on automation: industrial robots

The large amount of attention that ICT has received from researchers is well justified by its ubiquity. However, ICT encompasses a large number of specific technologies, including for instance word processing, databases, computer aided design, and the internet, with various degrees of overlap and interconnectedness. Therefore, the impact of ICT encompasses nearly the entire range of effects that technology in general may have on the labor market, including substitution of labor, complementarities with labor, organizational changes, and the creation of new tasks, among others. To better understand a specific aspect of technological change, one may prefer to study a narrower kind of technology. For instance, to investigate the economic impact of automation, researchers have focused on industrial robots, using the country-industry data on robot deliveries provided by the International Federation of Robotics (2019).

Combining the IFR data with country-industry data on value added, labor, and other capital from the EUKLEMS database, in Graetz and Michaels (2018) we document a strong positive association between robot adoption and productivity growth across 17 developed economies 1993-2007, as well as an association between robot adoption and a more intensive use of skilled labor, while not finding evidence for a negative impact on overall employment. In other words, industries that adopt robots at a faster pace do not appear to demand less labor in general, but appear to reduce their employment of low-skilled workers.

Acemoglu and Restrepo (forthcoming) instead study the impact of industrial robots across commuting zones in the US, using a location's initial employment distribution across industries to gauge exposure to robots. They find that exposure to robots locally is associated with reduced employment and lower wages, though the magnitudes are modest. In contrast, Dauth, Findeisen, Suedekum, and Woessner (2018) apply the same approach to German data and find that overall employment appears unaffected by increased robot use in a local labor market, although employment shifts from manufacturing to services. At the individual level, they find

that workers in exposed industries experience greater job stability but lower wage growth.

Finally, the impact of industrial robots has also been studied at the firm level. Koch, Manuylov, and Smolka (2019) find that Spanish manufacturing firms that adopted robots saw their labor cost share decrease, but their employment *increase*, relative to comparable firms who did not. Humlum (2019) shows similar evidence for Danish firms, where robot adoption was associated with an increase in sales, an increase in the wage bill, a reduction in the employment of production workers, and an increase in the employment of technicians and engineers.

From a theoretical point of view, this diversity of empirical findings is not surprising. As we show in Graetz and Michaels (2018), the impact of robot adoption on employment is theoretically ambiguous. Robot-adopting firms operate at lower cost, and are thus able to capture a larger share of the market. These firms will increase (lower, leave unchanged) their employment if the elasticity of demand they face is larger (smaller, the same) than the elasticity of substitution across tasks in the production process.

The impact of robots on wages is similarly ambiguous, since replacement and productivity effects work in opposite directions (Acemoglu and Restrepo, forthcoming). Moreover, these effects cannot be directly estimated empirically using variation across regions or firms, since such approaches difference out the general equilibrium adjustment of wages. Due to labor supply responses, the same problem inflicts the estimation of overall employment effects. Using structural models to take into account these general equilibrium adjustments, Acemoglu and Restrepo (forthcoming) estimate a negative impact of robots on employment and wages for the US economy, while Humlum (2019) estimates positive effects on both for the Danish economy.

In sum, the evidence on industrial robots demonstrates that even the relatively straightforward case of factory automation involves complex adjustment mechanisms already at the firm level, not to mention economy-wide general equilibrium responses. Effects on firm-level employment, economy-wide employment, or economy-wide wages are all theoretically ambiguous, and therefore it should not come as a surprise that different studies reach different conclusions. A systematic review of the evidence, and reconciliation of different findings in light of theory as well as institutional differences, has not been carried out yet. However, there is so

far no strong evidence that robotics-driven automation has caused massive job losses. On the contrary, robots' productivity effect may dominate their replacement effect, leading to higher employment and wages at the aggregate level, or even at the level of the individual firm. While Abraham and Kearney (2018) argue that robot adoption, after increased trade with China, is the second most important factor in accounting for the decline in the US employment-to-population ratio, this conclusion does not carry over to other OECD countries, where robots do not appear to reduce employment or wages, and where the employment rate has typically shown an upward trend in recent decades (OECD, 2019).

2.5 Summary

There is strong evidence that technological change—the ICT revolution in particular—has led to an increased demand for skilled labor, in the case of the US manifesting itself in a dramatic rise in the skill premium. There is similarly strong evidence that ICT has had an uneven impact across occupations, and has caused a polarization of occupational employment in most developed countries. Both findings are well accounted for by theories of skill biased and task-biased technological change, respectively.

Several researchers have argued that there has been a decline in the share of national income accruing to labor, as well as a rise in firms' market power, and that the two phenomena are related. However, others contest these claims, and there is as yet no consensus among economists. A fall in the labor share could also be explained by task-biased technological change, without reference to changing market power.

Finally, recent research has explored the effects of a narrowly defined automation technology, namely industrial robots. Outside the US, there is no evidence that robots have caused job losses, and even for the US the evidence points to only modest disemployment effects. Again, standard theory can accommodate these findings, as its prediction regarding the effects of automation on overall employment and wages are ambiguous, depending on the relative strengths of a variety of adjustment mechanisms.

3 Swedish labor market trends 1985-2017

This section examines trends in the Swedish labor market 1985-2017, including employment rates, the distribution of employment across education groups, sectors, and occupations, as well as wage inequality. My goal is to see if any of these trends shows obvious signs of having been affected by technological change, and in particular, to see whether technology might have been disrupting the labor market in ways that is potentially costly to workers. I will discuss each fact in light of the literature review of Section 2, and will provide some international comparisons.

The underlying data are the LISA database maintained by statistics Sweden, covering the population of Swedish residents aged 16-64 annually 1985-2017,¹¹ as well as the Wage Structure Statistics, covering the population of public sector workers and a large sample of private sector workers. I will also use measures of cognitive and non-cognitive ability from the military enlistment.

3.1 Trends in employment and composition

Given that automation has progressed steadily over the past four decades, an obvious question is how employment has evolved. Since the Swedish population continues to grow, the appropriate measure to examine is the employment-to-population ratio, also termed the employment rate. The relevant population is the working age population, people aged 16 to 64. Within this population, I consider two definitions of employment: a broad measure including those who were employed in November of a given year and earn no less than the base amount throughout the year;¹² and a full-time measure further requiring that annual earnings add up to at least eight times the monthly wage that is observed in November.

Figure 1 plots the overall and full-time employment rates separately by gender, and broken down further by age groups within each panel. The figure reveals stable employment rates for men and women in their thirties and forties, with the exception of increasing full-time employment of women in their forties. Employment rates among older men and women (aged

¹¹Later years contain older individuals, as well, but I restrict the analysis to ages 16-64 for consistency.

¹²The base amount is defined by law and used for administrative purposes by the Social Security Agency. It is SEK 47,300 in 2020.

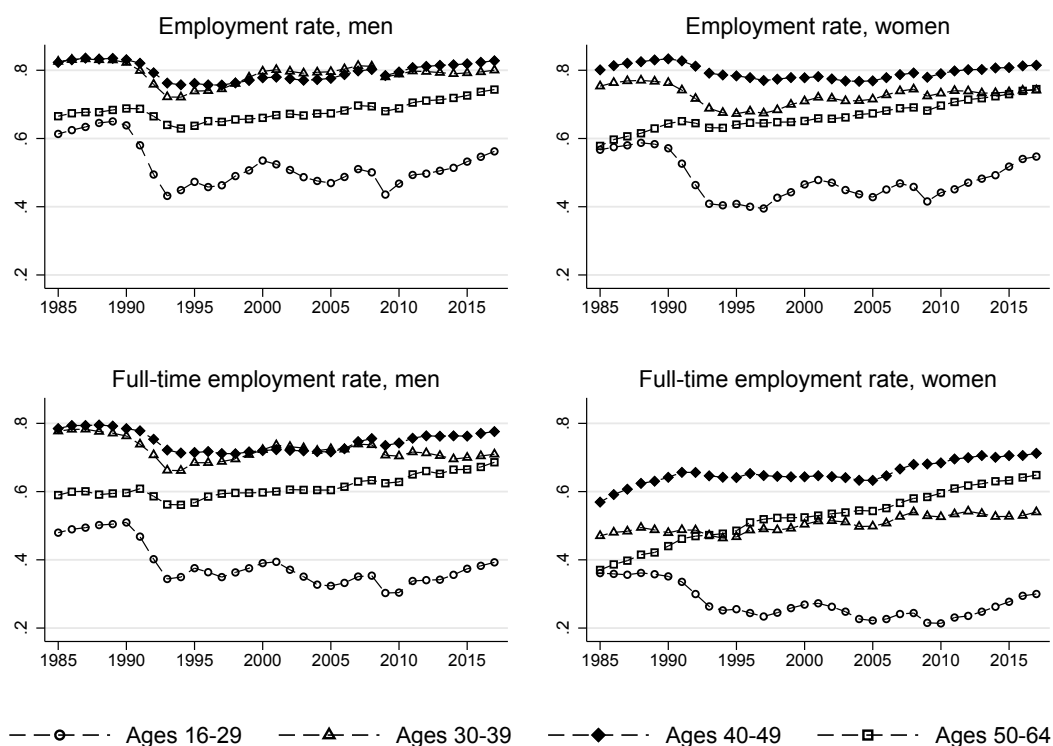


Figure 1: Employment rates across demographic groups

50-64) have been continuously increasing. Especially striking is the large increase in the full-time employment rate of older women, from about 40 percent in the late 1980s to about 65 percent in 2017. Finally, employment rates among the young (aged 16-30) were still somewhat lower in 2017 than at their peak prior to the early 1990s crisis. This is likely due to increased participation in tertiary education among this group.

The stable employment rates observed in Sweden are consistent with the absence of any strong evidence for disemployment effects of automation outside the US, as discussed in Section 2.4. Moreover, the majority of OECD countries have seen stable or increasing employment rates in recent decades, with the US being a prominent exception (OECD, 2019).

However, although overall employment rates appear stable in Sweden, the distribution of workers across sectors and occupations has undergone some dramatic changes. Figure 2 shows a sustained decline in the fraction employed in manufacturing of more than ten percentage points from a peak of nearly 25 percent. At the same time, the business services sector has increased from less than 10 to nearly 20 percent of total employment (this sector includes industries

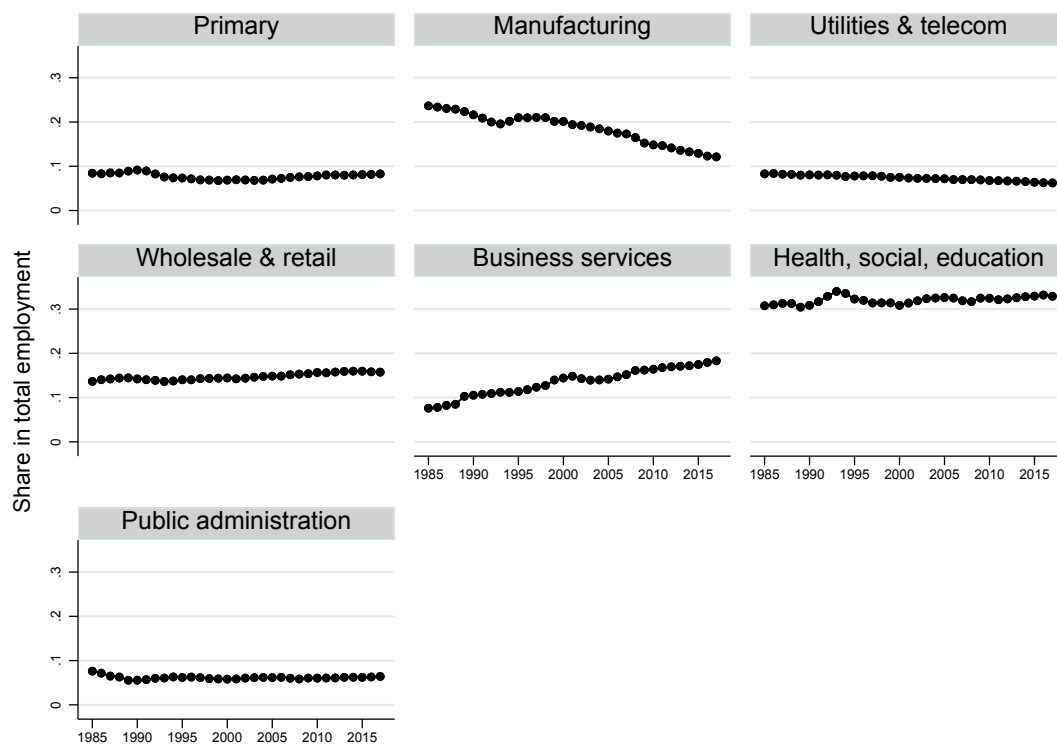


Figure 2: The distribution of employment across sectors

such as business and ICT consulting, insurance, and finance). A modest increase and decline, respectively, has been experienced by the utilities and sales sectors, while the remaining sectors appear stable.

Similarly, the occupational structure has changed substantially. Figure 3 plots employment shares of broad occupational groups 1996-2017.¹³ High-paying occupations such as ‘Officials & managers’ and Professionals’ employ a larger share of workers, while middle-paying white collar (‘Clerks’) and production-related (‘Crafts’, ‘Operators & assemblers’) occupations have declined. Figure 3 shows only mild upward trends for low-paying occupations (‘Service & sales’, ‘Elementary occupations’). However, a more detailed breakdown of occupations confirms that job polarization is occurring also in Sweden (Adermon and Gustavsson, 2015; Hensvik and Skans, 2019).

The Swedish experience of an employment shift from manufacturing to services, as well

¹³I obtain occupation data from the Wage Structure Statistics, where the occupation variable is available only since 1996. The subcategories of the broad occupations plotted in the figure are listed in Table A1.

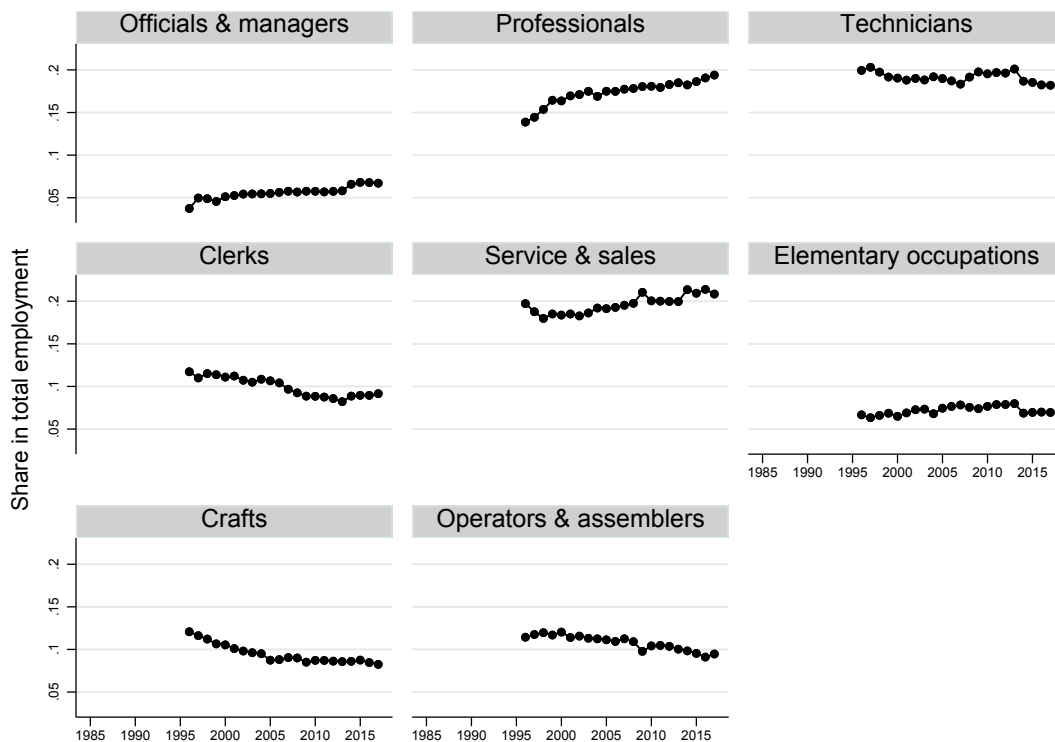


Figure 3: The distribution of employment across occupations

as job polarization, is well in line with trends in other countries. Job polarization is likely due to task-biased technological change (Goos, Manning, and Salomons, 2014)—declining occupations such as ‘Operators & assemblers’ and ‘Clerks’ used to be intensive in the use of repetitive and easily codifiable tasks which have been highly susceptible to automation (Autor, Levy, and Murnane, 2003).

The explanations for sectoral shifts are somewhat more involved. One potential explanation is differential productivity growth across sectors, leading to differential price growth, and causing consumers to increase the share of their expenditures devoted to the sector with relatively slow productivity growth.¹⁴ Another explanation is that consumers’ expenditures shift from ‘necessities’ to ‘luxuries’ as incomes rise, which would occur even if productivity growth was even across sectors. Empirically, both explanations receive support (Herrendorf, Rogerson, and Valentinyi, 2014). Given the substantial overlap between sectors and occupations (for instance,

¹⁴This mechanism requires that consumer demand is relatively inelastic across broad sectors. Given this, the mechanism does *not* require a particular form of technological change—it could substitute for labor or be complementary to it—as long as productivity growth is faster in manufacturing than services.

most operators and assemblers work in the manufacturing sector), there is also a case to be made for job polarization to be driven by structural change, at least in part (Barany and Siegel, 2018).

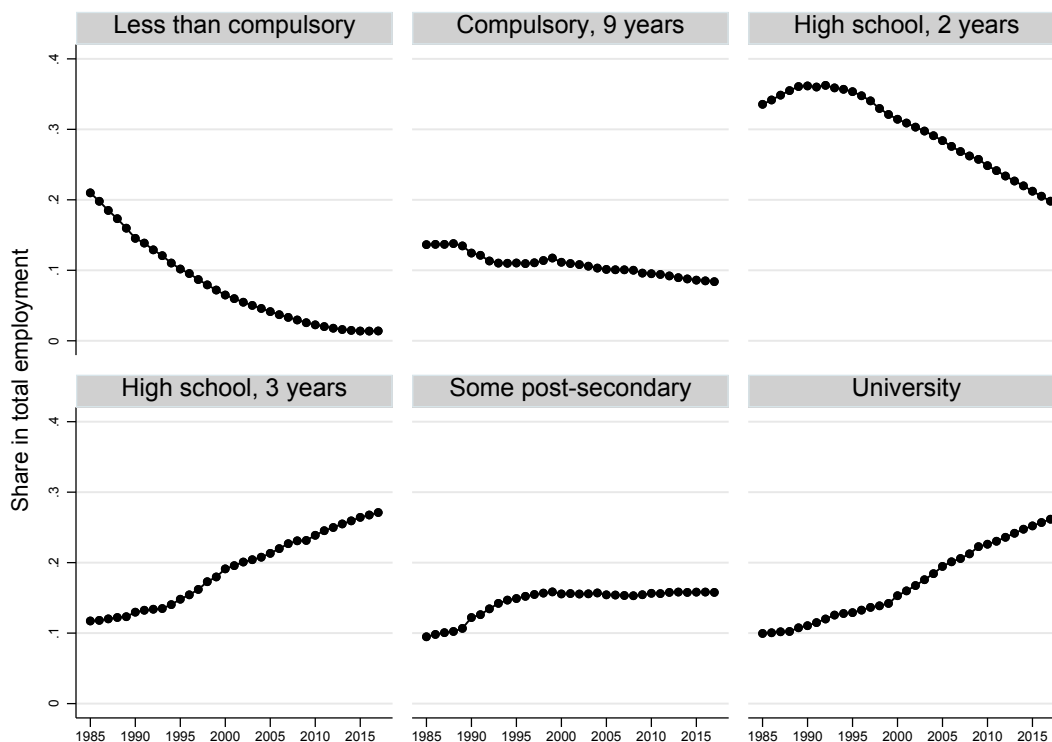


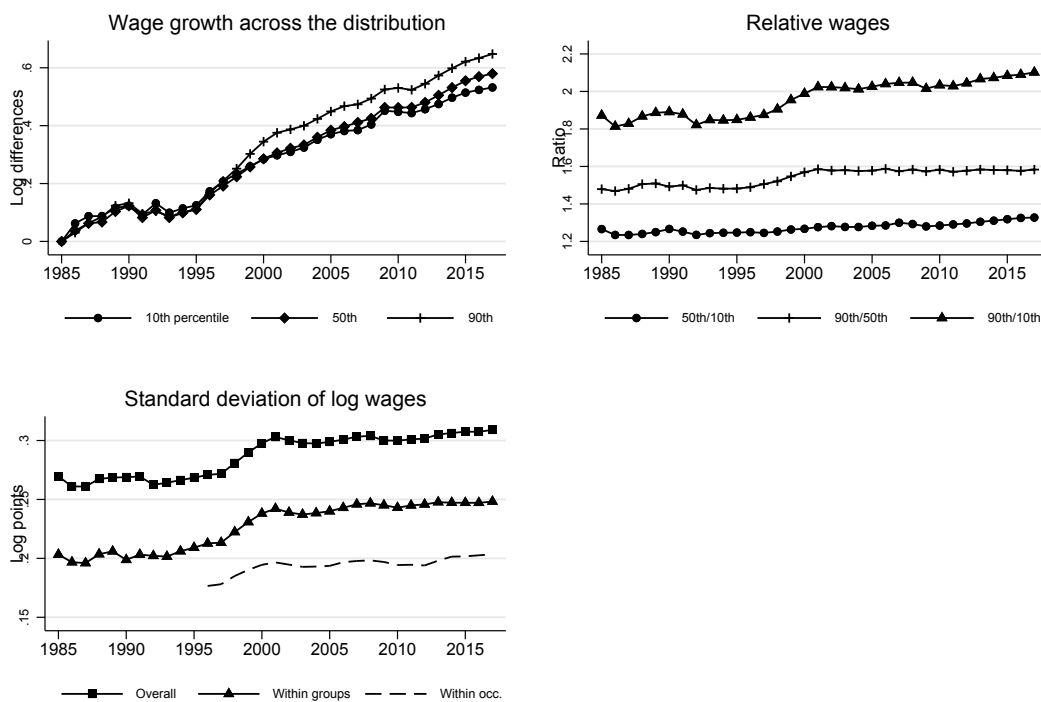
Figure 4: The educational composition of the workforce

Finally, I note large changes in the educational composition of the Swedish workforce. As seen in Figure 4, the fractions of university educated workers and those who completed a 3-year high school degree have grown substantially. Again, this experience is common across developed countries (OECD, 2020). The question arises how returns to education have evolved, given this large increase in supply. I address this question in the next sub-section.

3.2 Trends in wage inequality and skill returns

Here I examine trends in Swedish wage inequality and wage returns to education and skills. I focus on monthly wage rates throughout—the amount reported by the employer that a worker will receive if working full time for a month. I am interested in wage rates because they reflect the ‘price’ of labor more closely than labor earnings, which also depend on hours worked. For

the most part, labor earnings display similar trends as wages, and these results are available on request.¹⁵



Sample: Employees aged 16-64

Figure 5: Growth in real wages and trends in wage inequality

The top-left panel in Figure 5 plots the logarithm of real monthly wages for the median, the 10th percentile, and the 90th percentile from 1985 to 2017, as a difference to the initial (1985) level. Apart from the recessions in the early 1990s and early 2010s, there is steady wage growth throughout the distribution, but the three series start to diverge in the late 1990s. Another way to visualize this increased inequality is to plot ratios of the median to the 10th percentile, the 90th percentile to the median, and the 90th to the 10th percentile, as is done in the top-right panel. Inequality in the upper part of the distribution rises in the late 1990s but then flattens out, while bottom-half inequality rises mainly after 2000. Finally, the bottom-left panel of Figure 5 plots the standard deviation of log wages, showing a steep increase in the late 1990s and a much more modest rise since 2000 (squares).¹⁶ Wage inequality within demographic-education cells

¹⁵An exception is a larger increase in bottom-half inequality in terms of labor earnings than in terms of wages, especially among women.

¹⁶Edin and Holmlund (1995) document a strong decline in Swedish wage inequality throughout the 1970s and

shows a very similar trend (triangles), as does inequality within demographic-education cells and within occupations (dashed line).¹⁷

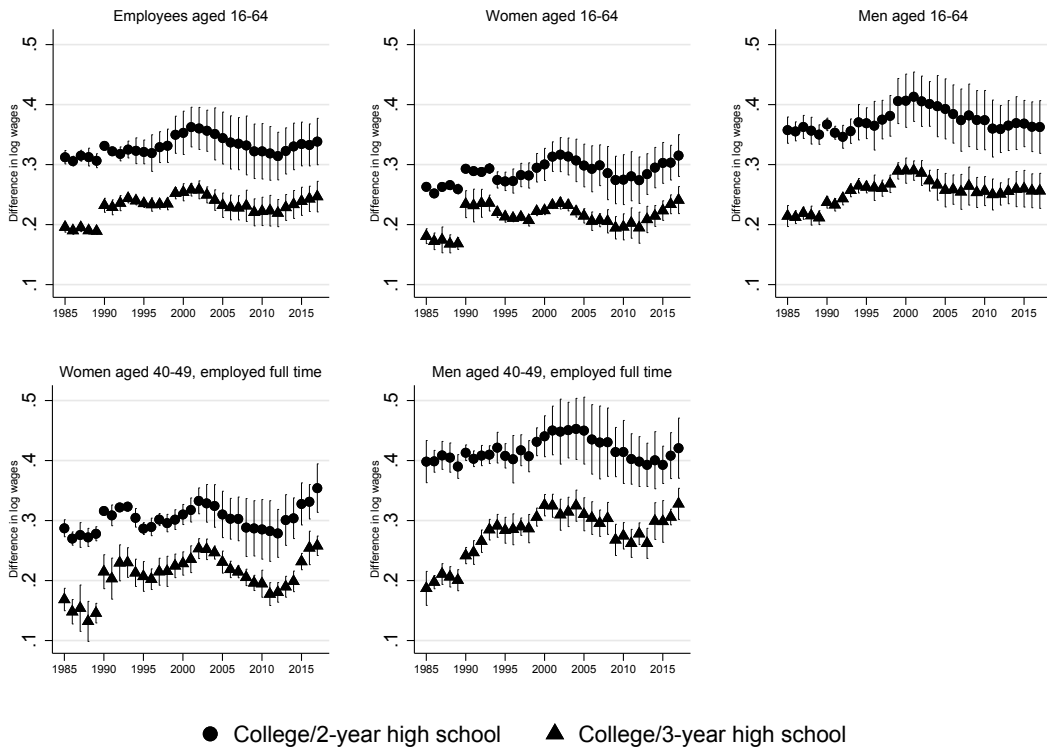
These wage inequality trends are similar when restricting attention to full-time workers (Figure A1), and they do not differ much between private and public sectors (Figures A2 and A3), but there are striking differences between men and women. Among men, top-half inequality has decreased somewhat since 2000, while bottom-half inequality has increased. Wage dispersion overall and within cells has stayed flat (Figure A4). In contrast, wage inequality among women has continued to increase throughout the distribution, overall and within groups (Figure A5). This difference holds also within the private and public sectors (Figures A6 and A7, A8 and A9), and is especially pronounced among prime-age workers (Figures A10 and A11). It also largely holds for within-occupation dispersion, meaning that increased representation of women in high-wage occupations (such as management) cannot be the sole explanation for these differential trends across the genders.

Next, I examine how returns to education and skills have evolved. I first regress wages on an exhaustive set of educational attainment dummies, along with a polynomial in potential experience and a dummies for gender and immigrant status. From this regression, I report the difference in log wages between college and a two-year high school education, as well as between college and a three-year high school degree. As shown in the top-left panel of Figure 6, the returns to college increased until the early 2000s, then decreased, and increased again in recent years. The returns have usually ranged from 30-40 log points (35-50 percent) for college versus two-year high school, and from 20-25 log points (22-28 percent) for college versus three-year high school. Again, there are differences across the genders, with the returns among women rising in recent years, while among men they stayed flat (top-right panels, bottom panels). Returns to college are still higher among men than among women, however.¹⁸

early 1980s.

¹⁷These measures are obtained as follows. First, I regress log wages on a fully interacted set of demographic (age, gender, immigrant status) and education dummies, and then calculate the standard deviation of the residuals ('within groups'). Second, I add to the same regression an exhaustive set of 3-digit occupation dummies (not interacted), and once more calculate the standard deviation of the residuals ('within occ.').

¹⁸Figure 6 shows striking jumps in skill returns from 1989 to 1990, especially so among women. I have verified that these jumps are not accounted for by spurious reclassification of educational attainment, or a discontinuity stemming from the sampling weights.



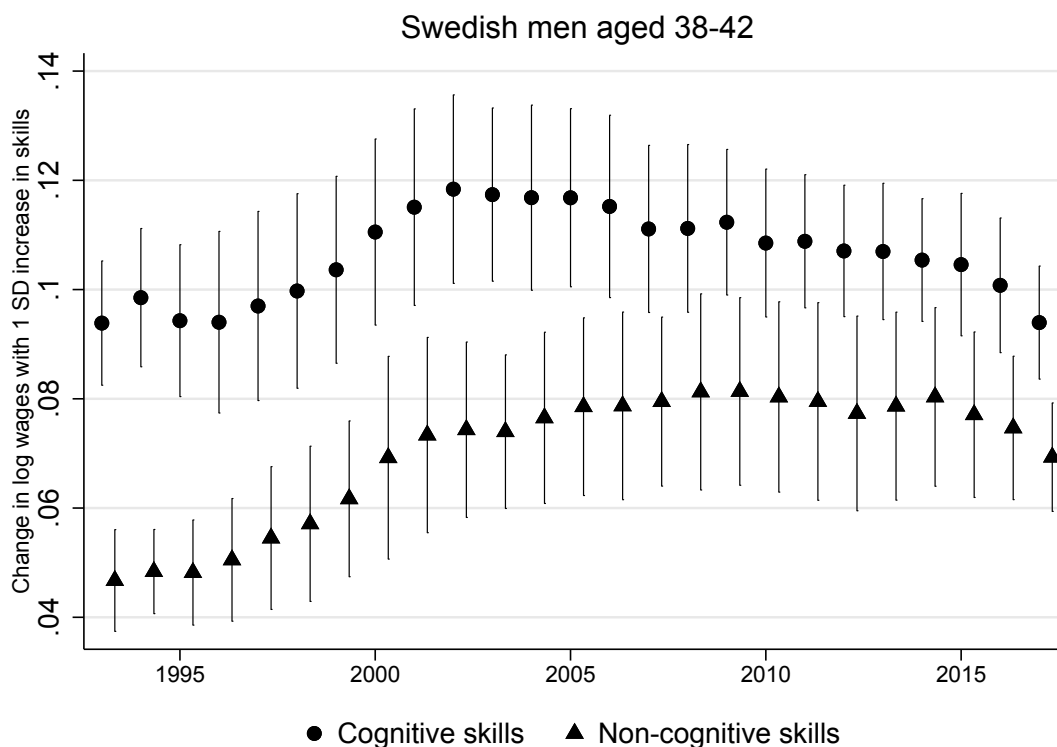
Notes: Markers and vertical lines represent point estimates and 95-percent confidence intervals, respectively. See the text for details of the estimation.

Figure 6: The college premium

Using data from military enlistment exams, it is also possible to examine wage returns to cognitive and non-cognitive (psycho-social) skills for men aged 38-42 over the period 1993-2017.¹⁹ Figure 7 confirms the findings of Edin, Fredriksson, Nybom, and Öckert (2017): a rise in the returns to cognitive skills until about 2000, but a larger and more sustained rise in the returns to non-cognitive skills until the late 2000s.

Unlike the facts related to employment described in Section 3.1, Swedish wage inequality trends are somewhat unusual relative to the experiences of some major developed economies. First, wage dispersion in Sweden is still much lower than elsewhere. For instance, the 90-10 ratio in Sweden stands at little over two (Figure 5), while in the US it is about five (Autor, Katz, and Kearney, 2008), in the UK it is about four (Butcher, Dickens, and Manning, 2012), and in France about three (Verdugo, Fraise, and Horny, 2012). Second, Swedish wage growth

¹⁹Extending the age range is possible for some years but not for the entire period, given that the enlistment data only cover the birth cohorts mid-1950s to early 1980s.



Notes: Markers and vertical lines represent point estimates and 95-percent confidence intervals, respectively. See the text for details of the estimation.

Figure 7: Returns to skill

has been more sustained than in the US and the UK, where real wage growth has stagnated or even turned negative over sustained periods (Acemoglu and Autor, 2011; Gregg, Machin, and Fernandez-Salgado, 2014). Third, there are differences in timing. In the US, top-half inequality continued to increase after 2000, while bottom-half dispersion actually fell (Acemoglu and Autor, 2011). This is the opposite of what happened in Sweden. In Germany, wage inequality continued to increase throughout the 2000s (Card, Heining, and Kline, 2013), while in Sweden it barely moved after 2000. Fourth, the skill premium has increased substantially in the US (Acemoglu and Autor, 2011), UK (Machin and Van Reenen, 1998; Lindley and Machin, 2016), and Germany (Dustmann, Ludsteck, and Schönberg, 2009), while it has not displayed an unambiguous trend in Sweden, and has increased only mildly over the entire period. However, when it comes to the returns to cognitive and psycho-social skills, trends in Sweden are similar to those in the US, as psycho-social skills have become more important in absolute and relative terms in both countries.

3.3 Summary and interpretation of employment and inequality trends

Despite an ever increasing scope for automation, there is no sign that jobs are harder to find in Sweden, as both employment rates and wage growth have remained stable. As argued in Section 2.4, this is consistent with economic theory, whose only unambiguous prediction about automation is that labor will be reallocated. Indeed, the occupational and sectoral distributions of employment have changed, and in ways similar to other developed countries. And as elsewhere, the educational attainment of the workforce has continued to increase.

It is the latter fact that, together with the stability of the college premium, also points to the presence of skill-biased technological change (SBTC) in the Swedish economy. In the absence of a demand shift favoring college-educated workers, an increase in skill supplies of the magnitude observed in Sweden (4) would have led to an unambiguous and sustained decline in the college premium. The question is, of course, why the college premium did not rise in Sweden as it has in other countries. But the absence of a strong decline nevertheless points to a skill-biased demand shift. Given the evidence from other countries discussed in Section 2.1, the most likely candidate for this demand shift is SBTC.

The presence of job polarization also suggests that task-biased technological change (TBTC) is operating in Sweden. However, the Swedish wage structure has not evolved as predicted by models of TBTC, in which job polarization is associated with widening dispersion at the top of the wage distribution but increased compression at the bottom (Costinot and Vogel, 2010; Feng and Graetz, forthcoming), the opposite of what happened since 2000. I will present some evidence in Section 4 that is consistent with TBTC. However, it is difficult to escape the conclusion that there are some salient aspects of changes in Swedish wage inequality that are unlikely to be related to technology, including the abrupt halt in the growth of inequality among men in 2000 and the continuing rise in wage dispersion among women. Institutional factors, including a time-varying degree of wage coordination in collective bargaining (Flam, 2019), as well as changing norms (Hsieh, Hurst, Jones, and Klenow, 2019) may play an important role, as well.

4 The occupation-level and regional impacts of technological change

In this section I summarize the results from two ongoing research projects that examine the uneven impact of technological change across occupations, and its consequence for the wage structure and workers' careers.

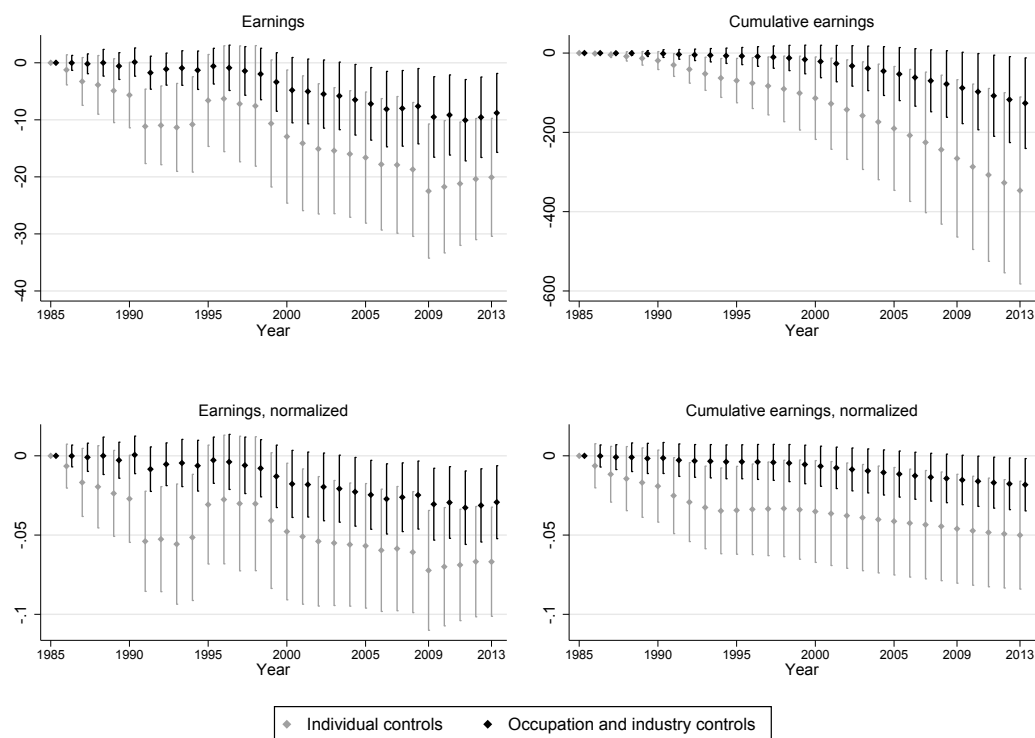
4.1 The career costs of occupational decline

In Edin, Evans, Graetz, Hernnäs, and Michaels (2019) we explore the career costs of occupational decline. We begin by identifying occupations that have declined sharply during the last 30 years and determine whether their decline was due to technological replacement using the Occupational Outlook Handbook (Bureau of Labor Statistics, 1986, 2017, OOH). We classify occupations as having declined if their employment in the US contracted by more than 25 percent. We then map this information to Swedish occupations in order to study how occupational decline affects individual workers, using data on the entire Swedish population at annual frequency 1985-2013.²⁰ We are also able to assess to what extent occupational decline was anticipated, using forecasts contained in the OOH, as well as the size and past growth of Swedish occupations (which strongly predict growth 1985-2013).

Although occupational decline represents a more gradual fall in demand compared to say a mass layoff (Jacobson, LaLonde, and Sullivan, 1993; Eliason and Storrie, 2006), we do find substantial costs for workers who in 1985 worked in a subsequently declining occupation. Over a period of 28 years, these workers have 2-5 percent lower cumulative earnings than comparable workers in non-declining jobs. And for workers at the bottom of the within-occupation earnings distribution, the losses are even larger at 8-11 percent. The range of estimates is based on several reasonable regression specifications: The upper end comes from comparing similar workers—in terms of gender, age, education, region, and prior income—across declining and non-declining

²⁰It is much more challenging to track employment changes for hundreds of different occupations in Sweden than in the US, because Swedish occupational classifications have changed substantially. However, at the level of detail at which we are able to consistently measure occupational employment in Sweden, we do see a strong association between our US-based indicator of decline and actual Swedish employment changes. Furthermore, studying the effects of occupational decline in the US is challenging given the lack of large longitudinal data sets. Nonetheless, we replicate our analysis using data from the National Longitudinal Survey of Youth. These results are much less precise than the Swedish ones, but lead to broadly similar conclusions.

occupations. The lower end is due to a more restrictive specification that compares similar workers in similar industries and occupations (including occupational employment forecasts).



Notes: Markers and vertical lines represent point estimates and 95-percent confidence intervals, respectively. See the text for details of the estimation.

Figure 8: Earnings costs of occupational decline


Figure 8 shows how the earnings losses of occupational decline accumulate over time. The top panel displays differences in terms of thousand SEK (2014 base), while the bottom panel displays differences in percentages of sample averages. The range of 2-5 percent of the mean mentioned above is represented by the rightmost black and grey markers in the bottom right panel, showing the differences in cumulative earnings at the end of our sample period, 2013. While the earnings losses for the most part build gradually over time, they appear particularly large during the recession in the early 1990s.

We also find that workers exposed to occupational decline are less likely to still be working in their initial occupation in 2013. This is noteworthy because over a nearly 30-year period, occupations could decline dramatically simply by taking in fewer younger workers and via regular retirements. Furthermore, we find that occupational decline is associated with increased

unemployment and publicly sponsored retraining. Our baseline results focus on all occupations that have declined, but we find very similar results when focusing on occupations whose decline was directly linked to technological change.

Occupational mobility is in principle a mechanism that may help workers mitigate their earnings losses from occupational decline. However, workers in declining occupations may also be more exposed to displacement, and given labor market frictions, may find themselves making occupational moves that are associated with higher earnings losses than incurred by those who manage to stay. We do not find that movers out of declining occupations do better than stayers in those same occupations. However, it is likely that a high rate of occupational mobility helps to reduce earnings losses because of general equilibrium effects, as it implies an upward-sloping occupational supply curve.

4.2 Task-biased technological change and occupational wage inequality



Notes: Employment and wage growth are measured relative to *clerks*. See the text for details of the estimation.

Figure 9: Occupational wage and employment growth

Workers' mobility response is part of the motivation for a second ongoing research project. In Adermon, Graetz, and Yakymovych (in progress), we explore whether changes in occupational wage premia are consistent with the notion that technological change has affected the demand for occupations differentially. Workers' mobility matters because it implies that changes in average occupational wages are affected by changes in composition. For instance, when a high-skill occupations such as professionals expands, the incoming workers may have lower productivity than the incumbents. This pushes down average wages, partly offsetting a demand-driven rise in the occupational 'price'—that is, the wage paid to a standardized unit of occupational output. When plotting raw occupational wage growth against employment growth in the left panel of Figure 9, we find no evidence of a positive relationship. Should this be seen as evidence against TBTC?

The task framework discussed in Section 2 predicts that in occupations where technology substitutes for labor, the occupational price should fall with technological progress, while it should rise in occupations where technology complements labor; and the former group of occupations should shrink, while the latter should grow. Thus, the task framework predicts a positive relationship between growth in occupational prices and employment growth. In contrast, the relationship between raw wage growth and employment changes is ambiguous, as wages are affected by compositional changes when workers entering or leaving an occupation are systematically different from those who stay. If entrants have lower productivity than incumbents—which will be the case if workers accumulate occupation-specific capital on the job—than the correlation between raw wage growth and employment changes will be smaller than that between price growth and employment changes.

To see whether the task framework is a good description of technology's impact, we would thus like to measure changes in occupational prices, which are different from changes in raw occupational wages due to workers' mobility response. We therefore pursue an alternative strategy, namely to compare the wage growth of stayers across occupations, that is, focussing on the within-occupation-spell wage growth of individual workers, following Cortes (2016). This eliminates the effect of a changing composition in terms of time-invariant abilities. And

after adjusting occupational wage growth in this way, we do indeed see a positive relationship with employment shifts—see the right panel of Figure 9. This evidence is consistent with the predictions from the task model: Workers shift to occupations that become more attractive over time—where the price of a unit of labor increases relatively more, as technology complements rather than substitutes for labor in these occupations; but entrants tend to have lower productivity than incumbents, so that occupational wage changes partly capture compositional shifts.²¹

There are however three concerns about the interpretation of our estimates. First, if workers' occupation-specific ability has a time-varying component, then stayers in declining occupations will be positively selected, whereas stayers in expanding occupations are negatively selected (assuming the causes for the decline and expansion are falling and rising occupational prices, respectively). This means that our estimates of changes in occupational prices will be attenuated. We have performed simulations suggesting that this bias is only of modest size. We are also exploring alternative estimation strategies which correct this bias that have been suggested in recent literature (Böhm, von Gaudecker, and Schran, 2019).

Second, skill accumulation over the life cycle likely differs across occupations, and this confounds our estimates. To address this concern, we restrict our sample to men aged 40-49, for whom skill accumulation is arguably rather modest regardless of occupation. This produces very similar results. Third, the changes in the returns to skills documented in Section 3.2 could vary by occupation, which would also confound our estimates. We are currently exploring this issue, and preliminary results suggest that skill returns have changed rather uniformly across occupation, which alleviates this concern.

Our findings are consistent with the conclusions of existing studies for the US (Cortes, 2016) and Germany (Böhm, von Gaudecker, and Schran, 2019), where wage and employment growth across occupations show a similarly weak correlation, but growth in occupational prices appears strongly positively related to growth in employment. Based on our findings, we are currently exploring to what extent TBTC has contributed to changes in wage inequality in Sweden.

²¹A direct comparison of entrants' and incumbents' wages confirms this interpretation. Results are available on request.

4.3 The impact of technological change across regions

In this section, I explore how the incidence of technological change varies across Swedish municipalities, and how local employment and wage bill growth correlate with this exposure. The technological changes I consider are the adoption of industrial robots and ICT over the period 1993-2007. I measure local exposure based on the distribution of employment across industries, following Acemoglu and Restrepo (forthcoming). More precisely, I first calculate the change in robot intensity (the number of robots per million hours worked) and the change in ICT intensity (the difference between the change in log ICT input and the change in log hours worked) for each of 28 two-digit industries. Second, in each municipality I sum over the product of an industry's municipal employment share (measured in 1990) and the industry's change in technology intensity.²²

Figure 10 shows how exposure to robots and ICT vary across Swedish municipalities, along with employment growth 1995-2017. Municipalities most exposed to robot adoption tended to be located in the southern half of the county and away from the big cities, while the geographical pattern of ICT exposure is somewhat difficult to characterize. Employment growth 1995-2017 was fastest in the large cities and surrounding areas, and slowest in rural municipalities.²³

Figure 11 plots employment growth against technology adoption across municipalities, along with population-weighted fitted lines. Employment growth is negatively correlated with robot adoption and positively correlated with ICT exposure.²⁴ But there are of course many other dimensions along which localities differ, which may be correlated both with technology adoption and employment growth. Similarly, adoption of different technologies may also be systematically correlated with each other. In general, any changes to product demand and labor

²²Data on robots (from the International Federation of Robotics) and ICT (from the EUKLEMS database) are available also for more recent years, but the coverage is less even. I obtain similar results if I compute regional outcomes over the period 1995-2007. I use 1990 industrial employment shares as Sweden had not recovered fully from the early 1990s recession in 1995, so that 1995 employment shares likely contain more transitory variation.

²³Municipalities most exposed to robot adoption include Trelleborg, Trollhättan, and Värnamo; least-exposed municipalities include Stockholm and its satellites Danderyd and Täby. Turning to ICT, most-exposed municipalities include Helsingborg, Sigtuna, and Karlshamn, while Karlskrona, Mariestad, and Motala are among the least-exposed.

²⁴When comparing Figures 10 and 11, one should keep in mind that many of the municipalities in Northern Sweden have small populations, and therefore do not much influence the population-weighted regression slopes in 11.

supply will affect both technology adoption and employment growth. To address some of these concerns, I explore three regression specifications as follows. First, I jointly include among the independent variables changes in robot use, ICT intensity, and non-ICT capital, while controlling for initial employment and the employment-to-population ratio (in logs). Second, I control in addition for the educational composition in each region (in terms of the three groups less than high school, high school, and college plus). Third, I additionally control for the industrial employment mix (in terms of the four broad sectors manufacturing, services, utilities, and primary).

Panel A in Table 1 displays the results for employment growth. The raw correlations between employment growth and technological changes shown in Figure 11 remain qualitatively unchanged when entering the technology variables jointly and controlling for initial employment and population (column (1)) and when controlling for skill mix (column (2)). However, when also controlling for industrial composition, employment growth is positively related to both robot and ICT adoption (column (3)). The magnitudes can be interpreted as follows. If the change in robots per million hours worked (about 500 employees working full-year-full-time) in a municipality had been zero instead of its actual average of 0.31, employment growth would have been 4 percentage points lower ($100 \times 0.31 \times 0.13$). Or, if ICT intensity had changed by ten percent less, employment growth would have been 3.8 percentage points lower ($100 \times 0.1 \times 0.38$). One must keep in mind, of course, that these results represent conditional correlations and do not imply a causal interpretation. Nevertheless, they go some way in informing the question whether technological changes, including factory automation, are associated with job losses across regions.²⁵

Using the same regression specifications, I have explored a variety of further outcomes. Panel B of Table 1 shows that the relationship between growth in the employment-to-population ratio and technological change across regions is rather weak, although these results are suggestive of a positive association with increased ICT intensity. Panel C indicates that growth in municipality-level wage bills is similarly positively related to technology adoption as is em-

²⁵Prior research often employs instrumental variables—in particular, technology adoption in other countries—to get closer to a causal interpretation. Such an approach turned out to be infeasible in my setting, as estimates were much too imprecise to be informative.

Table 1: Technological change across regions: Employment and wage bill growth

	(1)	(2)	(3)
<i>A. Employment growth 1995-2017 (weighted mean: 0.23)</i>			
Robots	-0.028 (0.031)	-0.019 (0.033)	0.13 (0.038)
ICT	0.66 (0.16)	0.70 (0.16)	0.38 (0.17)
<i>B. Growth in employment-to-population ratio 1995-2017 (weighted mean: 0.13)</i>			
Robots	-0.016 (0.014)	-0.010 (0.014)	0.0087 (0.017)
ICT	0.14 (0.064)	0.13 (0.059)	0.058 (0.065)
<i>C. Growth in wage bill 1995-2017 (weighted mean: 0.67)</i>			
Robots	-0.071 (0.033)	-0.046 (0.034)	0.13 (0.041)
ICT	0.76 (0.17)	0.85 (0.17)	0.48 (0.18)
Other capital	✓	✓	✓
Employment and population controls	✓	✓	✓
Education controls		✓	✓
Industry controls			✓

Notes: Results are from regressions of the dependent variables indicated in the panel headings on robot adoption, the change in ICT intensity, and controls. The sample consists of 284 Swedish municipalities. ‘Other capital’ indicates that the change in non-ICT intensity is controlled for. Employment and population controls include the log of the employment-to-population ratio as well as the log of employment in 1990. Education controls include the employment shares of three educational groups, and industry controls include the employment shares of four broad sectoral groups, both in 1990. Regressions are weighted by initial employment shares. Robust standard errors in parentheses.

ployment growth.

In Table 2, I explore how technological change is related to changes in skill mix. While there is no clear pattern regarding robot adoption, increased ICT intensity is strongly associated with growth in the employment share of skilled workers. Moreover, it is more negatively related to growth in middle- than to growth in low-skill shares, which is at least somewhat suggestive of skill polarization across regions. These patterns are very similar when instead using skill shares in the wage bill as dependent variables (Table A2). Thus, regional patterns of technology adoption and skill demand in Sweden are in line with country-industry evidence when it comes to ICT (Michaels, Natraj, and Van Reenen, 2014) but not when it comes to robots (Graetz and Michaels, 2018).

There is some evidence that robot adoption is associated with a reallocation of employment from manufacturing to services, as in Dauth, Findeisen, Suedekum, and Woessner (2018), but this association is not very robust (Table A3).

In sum, the relationship between technology adoption and growth in employment and wage income across Swedish municipalities tends to be positive—there is no evidence that fast-adopting locations shed workers. On the other hand, there is some indication that new technologies—ICT in particular—are associated with higher inequality, as skilled workers appear to be favored. I must note, however, that the simple exercises I have performed here cannot capture several more complex ways in which technology may affect regional inequality. For instance, ICT likely affected the nature of agglomeration (Michaels, Rauch, and Redding, 2018), and the presence of high-tech industries likely implies spill-overs to the demand for low-skilled services (Moretti, 2010). To what extent these effects are present in Sweden is a matter for future research.

5 The future

In the face of impressive technological advances such as machine learning and robotics, policy makers and the public are understandably concerned about the adverse consequences of

Table 2: Technological change across regions: Skill share changes (employment)

	(1)	(2)	(3)
<i>A. Change in high-skill employment share 1995-2017 (weighted mean: 0.13)</i>			
Robots	-0.0018 (0.0093)	-0.0013 (0.0091)	-0.016 (0.011)
ICT	0.095 (0.046)	0.11 (0.044)	0.13 (0.044)
<i>A. Change in middle-skill employment share 1995-2017 (weighted mean: -0.02)</i>			
Robots	0.028 (0.0098)	0.0027 (0.0073)	-0.010 (0.0095)
ICT	-0.084 (0.050)	-0.11 (0.040)	-0.080 (0.039)
<i>A. Change in low-skill employment share 1995-2017 (weighted mean: -0.11)</i>			
Robots	-0.025 (0.0088)	0.000064 (0.0070)	0.024 (0.0092)
ICT	-0.020 (0.039)	-0.0077 (0.020)	-0.057 (0.021)
Other capital	✓	✓	✓
Employment and population controls	✓	✓	✓
Education controls		✓	✓
Industry controls			✓

Notes: Results are from regressions of the dependent variables indicated in the panel headings on robot adoption, the change in ICT intensity, and controls. The sample consists of 284 Swedish municipalities. ‘Other capital’ indicates that the change in non-ICT intensity is controlled for. Employment and population controls include the log of the employment-to-population ratio as well as the log of employment in 1990. Education controls include the employment shares of three educational groups, and industry controls include the employment shares of four broad sectoral groups, both in 1990. Regressions are weighted by initial employment shares. Robust standard errors in parentheses.

automation.²⁶ However, as noted in the introduction, experts disagree on the expected labor market impact of new technologies over the next few decades. Among those who expect rapid technological progress, some worry about the implications for the average worker (Brynjolfsson and McAfee, 2014; Ford, 2015), while others appear more optimistic (Autor, 2015). And there are those who are not even confident in the ability of machine learning and robotics to deliver sustained productivity growth (Gordon, 2012, 2014). In 2017, a panel of expert economists were asked to evaluate the statement “Holding labor market institutions and job training fixed, rising use of robots and artificial intelligence is likely to increase substantially the number of workers in advanced countries who are unemployed for long periods”. 38 percent agreed, 33 disagreed, and 29 percent were uncertain.²⁷

It should be noted that forecasting occupational employment trends has a long tradition in economic and policy research. Policy makers have sought to provide detailed and up-to-date projections of employment growth, as well as guidance about what types of skills can be expected to remain in demand. Examples include the forecasts by the Bureau of Labor Statistics (BLS) discussed above, the O*NET project sponsored by the US Department of Labor, and the Skills Forecast conducted by the European Centre for the Development of Vocational Training. The BLS forecasts have proved quite accurate in forecasting occupational employment trends not only in the US (Veneri, 1997) but also in Sweden (Edin, Evans, Graetz, Hernnäs, and Michaels, 2019). This is testament both to the quality of the forecasts as well as to the fact that labor demand shifts follow very similar patterns across countries.

While these projections have great merits, they do not directly engage with the characteristics of new technologies, and they cannot address claims that “This time is different.” A rapidly growing literature develops predictions about the future of work that do attempt to take into account the capabilities of new technologies. For the most part, this literature focuses on the expected labor market impact of machine learning (ML), widely regarded as the dominant paradigm in artificial intelligence research since the early 2010s (Mitchell, 2019; Somers, 2017).

²⁶This paragraph and the next draw heavily on Graetz (2019).

²⁷See <http://www.igmchicago.org/surveys/robots-and-artificial-intelligence>.

Agrawal, Gans, and Goldfarb (2019) argue that ML does not represent progress towards *general* artificial intelligence, but rather marks an advance in one specific area, namely *prediction*. However, prediction tasks are a very common input to decision tasks, and as such ML has wide applicability, making it general purpose technology such as the steam engine, the personal computer, or the internet. Agrawal, Gans, and Goldfarb (2019) then argue that the effects of ML will involve automation of prediction tasks such as mortgage approval (where one tries to predict whether the borrower will default) or hiring (predicting which applicant will perform best). As in the case of already established technologies, and as formalized by the task model discussed in Section 2.2, the impact on the demand for human labor then depends on how the negative effect of machine substitution compares to its productivity-enhancing effects—since prediction is complementary to decision tasks, some of which will still be carried out by humans—as well as the rate at which new types of decision tasks are created. Employing similarly qualitative reasoning, Brynjolfsson and Mitchell (2017) argue that ML will be used in tasks that map well-defined inputs to well-defined outputs, for which data sets exist linking input-output pairs, when there is some tolerance for error, and the function being learned does not change rapidly over time, among other things.

As such qualitative reasoning leads to ambiguous conclusions about the expected impact of ML on labor demand, other studies take a quantitative approach. Frey and Osborne (2017, FO) assign automation probabilities to the universe of occupations in the US, using a method that involves a survey of ML experts, identification of engineering bottlenecks, and an algorithm to extrapolate from a set of hand-labelled occupations. They find that the automation probability declines monotonically both in occupational wages as well as education, and predict that workers in transportation and logistics occupations, office and administrative support workers, production workers, as well as sales workers are at the highest risk of automation. In contrast, Webb (2020) estimates higher automation probabilities for high-skill occupations than for middle- and low-skill ones, based on the overlap of job task descriptions and ML patent texts.²⁸ Arntz, Gregory, and Zierahn (2017) argue that Frey and Osborne (2017) overestimate overall

²⁸Mann and Püttmann (2017) also use patent text analysis to measure automation exposure, but to shed light on the impact that automation has already had, rather than to predict its future impact.

automation risk, and find that only 9 percent of US employment is at risk of automation once within-occupation heterogeneity in tasks is taken into account, unlike the 47 percent estimated by Frey and Osborne (2017).

In the context of the Swedish labor market, the availability of detailed skills data from the military enlistment allows for a deeper investigation of the expected skill bias of future technologies. Hensvik and Skans (2019) find a positive relationship between the average skills of an occupation's workforce and the occupation's projected employment growth, both in terms of the Frey and Osborne (2017) measure as well as the more traditional employment projections calculated by the US Bureau of Labor Statistics. While this correlation holds for a composite skill measure, the multi-dimensionality of skills does matter. Occupations intensive in social maturity, emotional intensity, and verbal skills are projected to grow, while occupations intensive in inductive skills and emotional stability are projected to decline. Moreover, these correlations are largely unchanged when instead focusing on *actual* employment growth since 2001. Thus, recent changes in the demand for skills are expected to persist.

In a similar spirit, I explore how automation risk as measured by FO varies across municipalities, and how it is associated with recent employment growth at the regional level. To calculate regional automation risk, I assign each worker the FO measure based on her 3-digit occupation in 2013, and then simply take the average across workers in each municipality.²⁹ Figure 10 reveals that automation risk tends to be highest in the more rural parts of Southern Sweden—it is highest in Karlstad, Mariestad, and Värnamo, and lowest in Danderyd, Linköping, and Uppsala. And Figure 11 indicates that automation risk tends to be higher in municipalities that saw slower growth in the past decades. This suggests that adjustments to technological change required in the future are perhaps not much different from the adjustments that are already ongoing—employment is already shifting away from areas expected to be more affected by future automation. To explore this further, I run regressions similar to the ones reported in Section 4.3. Indeed, the negative relationship between recent employment growth and future automation risk is affected by controls. In particular, it is reduced to about 15 percent of its

²⁹I use 2013 to calculate projected automation as in 2014 a new occupational classification was introduced that is more difficult to match to the FO data.

original strength when comparing automation risk across municipalities with similar size, skill mix, and industry composition (Table A4). This means that automation risk is particularly high in municipalities with a large share of low-skill employment or a large manufacturing share, municipalities which have already been on a downward trajectory.

Thus, even the kinds of prediction exercises that attempt to explicitly take into account the capabilities of new technologies, tend to produce projections that are strongly related to recent sectoral and occupational shifts, indicating a high degree of persistence in the skill-bias and regional incidence of technological change (an exception is Webb, 2020). Of course, these projections are highly uncertain. One extreme of the range of possible outcomes is the arrival of human-level AI within a few decades. Indeed, AI experts commonly predict this to occur within 15-25 years from the date the prediction is made (Armstrong and Sotala, 2015), and have done so since the 1960s (Mitchell, 2019). The other extreme is imminent stagnation, as ML may soon be hitting drastically diminishing returns (Somers, 2017). Given this uncertainty, the best that researchers can do, arguably, is to continue the study of new technologies in terms of their scope for task replacement, complementarities, and the creation of new tasks, and to combine this understanding with an up-to-date description of occupational task content and skill requirements.

6 Conclusions

During the past three decades the Swedish labor market experienced the kinds of transformations one would expect in a modern and innovative economy, without suffering obvious disruptions, apart from the cyclical downturns that are unrelated to secular technological changes. Labor force participation remains high, and among older individuals has been steadily increasing. Wage growth has been steady. Inequality remains very low by international standards, although it has increased somewhat.

While technological change has plausibly affected labor demand in Sweden in similar ways as in other developed countries, these effects are more muted and somewhat hidden, likely because of Swedish wage setting institutions. The college premium in Sweden has remained

largely unchanged, in contrast to the near doubling in the US; but the absence of a marked *decline* in the premium nevertheless points to an increased demand for skills, given the large increase in the supply of college-educated labor that has occurred at the same time. An increase in the demand for skills is also clearly visible in the returns to cognitive and especially psychosocial skills, though it is puzzling why this trend has flattened out.

The fact that workers affected by occupational decline do suffer substantial earnings losses over the course of their careers is further evidence of the distributional impact of technological change—at the same time, the increased participation in public training by these workers is testament to a strong safety net. And occupational wage rates, once adjusted for compositional changes, do evolve in ways consistent with a task-biased impact of technology. However, some salient features of Swedish inequality are unlikely to be explained by technological change, especially the abrupt halt in the growth of inequality among men in 2000 and the continuing rise in wage dispersion among women.³⁰

Systematic projections of the future impact of technological change indicate a substantial degree of persistence. Skill-intensive occupations are expected to continue to grow disproportionately, as are regions with a high share of skilled employment. While these projections are of course uncertain, Swedish labor market institutions appear to be in a good position to deal with the disruptive effects of technological change.

³⁰Beaudry and Lewis (2014) present evidence for a strong link between the decline in the gender wage gap and the ICT revolution in the US, as the skills that are complementary to ICT appear to be more abundant among women. However, it is not clear how their explanation would apply to differential changes in wage dispersion across genders, or how it could account for the timing of the divergence in Sweden.

References

- ABRAHAM, K. G., AND M. S. KEARNEY (2018): “Explaining the Decline in the U.S. Employment-to-Population Ratio: A Review of the Evidence,” Working Paper 24333, National Bureau of Economic Research.
- ACEMOGLU, D. (2002): “Technical Change, Inequality, and the Labor Market,” *Journal of Economic Literature*, 40(1), 7–72.
- ACEMOGLU, D., AND D. AUTOR (2011): “Skills, Tasks and Technologies: Implications for Employment and Earnings,” in *Handbook of Labor Economics*, ed. by O. Ashenfelter, and D. Card, vol. 4, chap. 12, pp. 1043–1171. Elsevier.
- ACEMOGLU, D., AND P. RESTREPO (2018a): “Modeling Automation,” *AEA Papers and Proceedings*, 108, 48–53.
- (2018b): “The Race between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment,” *American Economic Review*, 108(6).
- (2019a): “Automation and New Tasks: How Technology Displaces and Reinstates Labor,” *Journal of Economic Perspectives*, 33(2), 3–30.
- (2019b): “The Wrong Kind of AI? Artificial Intelligence and the Future of Labor Demand,” Working Paper 25682, National Bureau of Economic Research.
- (forthcoming): “Robots and Jobs: Evidence from US Labor Markets,” *Journal of Political Economy*.
- ADERMON, A., G. GRAETZ, AND Y. YAKYMOVYCH (in progress): “Wage Growth across Occupations, Industries, and Regions,” mimeo.
- ADERMON, A., AND M. GUSTAVSSON (2015): “Job Polarization and Task-Biased Technological Change: Evidence from Sweden, 1975–2005,” *Scandinavian Journal of Economics*, 117(3), 878–917.

- AGRAWAL, A., J. S. GANS, AND A. GOLDFARB (2019): “Artificial Intelligence: The Ambiguous Labor Market Impact of Automating Prediction,” *Journal of Economic Perspectives*, 33(2), 31–50.
- AKERMAN, A., I. GAARDER, AND M. MOGSTAD (2015): “The Skill Complementarity of Broadband Internet *,” *The Quarterly Journal of Economics*, 130(4), 1781–1824.
- ARMSTRONG, S., AND K. SOTALA (2015): “How we’re predicting AI—or failing to,” in *Beyond artificial intelligence*, pp. 11–29. Springer.
- ARNTZ, M., T. GREGORY, AND U. ZIERAHN (2017): “Revisiting the risk of automation,” *Economics Letters*, 159, 157 – 160.
- ATALAY, E., P. PHONGTHIENGTHAM, S. SOTELO, AND D. TANNENBAUM (forthcoming): “The Evolution of Work in the United States,” *American Economic Journal: Applied*.
- AUM, S., D. KOH, AND R. SANTAELULIA-LLOPIS (2019): “Increasing and Decreasing Labor Shares: Cross-Country Differences in the XXI Century,” Working Papers 1135, Barcelona Graduate School of Economics.
- AUTOR, D., D. DORN, L. F. KATZ, C. PATTERSON, AND J. VAN REENEN (forthcoming): “The Fall of the Labor Share and the Rise of Superstar Firms,” *The Quarterly Journal of Economics*.
- AUTOR, D., C. GOLDIN, AND L. F. KATZ (2020): “Extending the Race between Education and Technology,” Working Paper 26705, National Bureau of Economic Research.
- AUTOR, D. H. (2015): “Why Are There Still So Many Jobs? The History and Future of Workplace Automation,” *Journal of Economic Perspectives*, 29(3), 3–30.
- (2019): “Work of the Past, Work of the Future,” *AEA Papers and Proceedings*, 109, 1–32.
- AUTOR, D. H., L. F. KATZ, AND M. S. KEARNEY (2006): “The Polarization of the U.S. Labor Market,” *American Economic Review*, 96(2), 189–194.

- (2008): “Trends in U.S. Wage Inequality: Revising the Revisionists,” *The Review of Economics and Statistics*, 90(2), 300–323.
- AUTOR, D. H., L. F. KATZ, AND A. B. KRUEGER (1998): “Computing Inequality: Have Computers Changed the Labor Market?,” *The Quarterly Journal of Economics*, 113(4), 1169–1213.
- AUTOR, D. H., F. LEVY, AND R. J. MURNANE (2003): “The Skill Content of Recent Technological Change: An Empirical Exploration*,” *The Quarterly Journal of Economics*, 118(4), 1279–1333.
- AZAR, J. A., I. MARINESCU, M. I. STEINBAUM, AND B. TASKA (2018): “Concentration in US Labor Markets: Evidence From Online Vacancy Data,” Working Paper 24395, National Bureau of Economic Research.
- BARANY, Z. L., AND C. SIEGEL (2018): “Job Polarization and Structural Change,” *American Economic Journal: Macroeconomics*, 10(1), 57–89.
- BARKAI, S. (2017): “Declining Labor and Capital Shares,” mimeo.
- BASU, S. (2019): “Are Price-Cost Markups Rising in the United States? A Discussion of the Evidence,” *Journal of Economic Perspectives*, 33(3), 3–22.
- BAZIKI, S. B. (2015): *Firms, International Competition, and the Labor Market*. Doctoral dissertation, Uppsala University.
- BAZIKI, S. B., R. GINJA, AND T. B. MILICEVIC (2018): “Trade Competition, Technology, and Labor Reallocation,” mimeo.
- BEAUDRY, P., D. A. GREEN, AND B. M. SAND (2016): “The Great Reversal in the Demand for Skill and Cognitive Tasks,” *Journal of Labor Economics*, 34(S1), 199–247.
- BEAUDRY, P., AND E. LEWIS (2014): “Do Male-Female Wage Differentials Reflect Differences in the Return to Skill? Cross-City Evidence from 1980-2000,” *American Economic Journal: Applied Economics*, 6(2), 178–194.

- BERMAN, E., J. BOUND, AND S. MACHIN (1998): “Implications of Skill-Biased Technological Change: International Evidence,” *The Quarterly Journal of Economics*, 113(4), 1245–1279.
- BÖHM, M. J., H.-M. VON GAUDECKER, AND F. SCHRAN (2019): “Occupation Growth, Skill Prices, and Wage Inequality,” mimeo.
- BRYNJOLFSSON, E., AND A. MCAFEE (2014): *The second machine age: Work, progress, and prosperity in a time of brilliant technologies*. WW Norton & Company.
- BRYNJOLFSSON, E., AND T. MITCHELL (2017): “What can machine learning do? Workforce implications,” *Science*, 358(6370), 1530–1534.
- BUREAU OF LABOR STATISTICS (1986): *Occupational Outlook Handbook, 1986-87 Edition : Bulletin of the United States Bureau of Labor Statistics, No. 2250* United States Department of Labor.
- (2017): *Occupational Outlook Handbook* United States Department of Labor.
- BUTCHER, T., R. DICKENS, AND A. MANNING (2012): “Minimum Wages and Wage Inequality: Some Theory and an Application to the UK,” CEP Discussion Papers dp1177, Centre for Economic Performance, LSE.
- CARD, D., AND J. E. DINARDO (2002): “Skill-Biased Technological Change and Rising Wage Inequality: Some Problems and Puzzles,” *Journal of Labor Economics*, 20(4), 733–783.
- CARD, D., J. HEINING, AND P. KLINE (2013): “Workplace Heterogeneity and the Rise of West German Wage Inequality,” *The Quarterly Journal of Economics*, 128(3), 967–1015.
- CASELLI, F., AND A. MANNING (forthcoming): “Robot Arithmetic: New Technology and Wages,” *American Economic Review: Insights*.
- CORTES, G. M. (2016): “Where Have the Middle-Wage Workers Gone? A Study of Polarization Using Panel Data,” *Journal of Labor Economics*, 34(1), 63–105.

- COSTINOT, A., AND J. VOGEL (2010): “Matching and Inequality in the World Economy,” *Journal of Political Economy*, 118(4), 747–786.
- DAUTH, W., S. FINDEISEN, J. SUEDEKUM, AND N. WOESSNER (2018): “Adjusting to Robots: Worker-Level Evidence,” Working Papers 13, Federal Reserve Bank of Minneapolis, Opportunity and Inclusive Growth Institute.
- DE LOECKER, J., AND J. EECKHOUT (2018): “Global Market Power,” Working Paper 24768, National Bureau of Economic Research.
- DE LOECKER, J., J. EECKHOUT, AND G. UNGER (forthcoming): “The Rise of Market Power and the Macroeconomic Implications,” *The Quarterly Journal of Economics*.
- DEMING, D. J. (2017): “The Growing Importance of Social Skills in the Labor Market*,” *The Quarterly Journal of Economics*, 132(4), 1593–1640.
- DiNARDO, J., N. M. FORTIN, AND T. LEMIEUX (1996): “Labor Market Institutions and the Distribution of Wages, 1973-1992: A Semiparametric Approach,” *Econometrica*, 64(5), 1001–1044.
- DiNARDO, J. E., AND J.-S. PISCHKE (1997): “The Returns to Computer Use Revisited: Have Pencils Changed the Wage Structure Too?,” *The Quarterly Journal of Economics*, 112(1), 291–303.
- DUSTMANN, C., J. LUDSTECK, AND U. SCHÖNBERG (2009): “Revisiting the German Wage Structure,” *The Quarterly Journal of Economics*, 124(2), 843–881.
- EDIN, P.-A., T. EVANS, G. GRAETZ, S. HERNNÄS, AND G. MICHAELS (2019): “Individual Consequences of Occupational Decline,” working paper.
- EDIN, P.-A., P. FREDRIKSSON, M. NYBOM, AND B. ÖCKERT (2017): “The Rising Return to Non-Cognitive Skill,” IZA Discussion Papers 10914, Institute for the Study of Labor (IZA).
- EDIN, P.-A., AND B. HOLMLUND (1995): *The Swedish Wage Structure: The Rise and Fall of Solidarity Wage Policy?* pp. 307–344. University of Chicago Press.

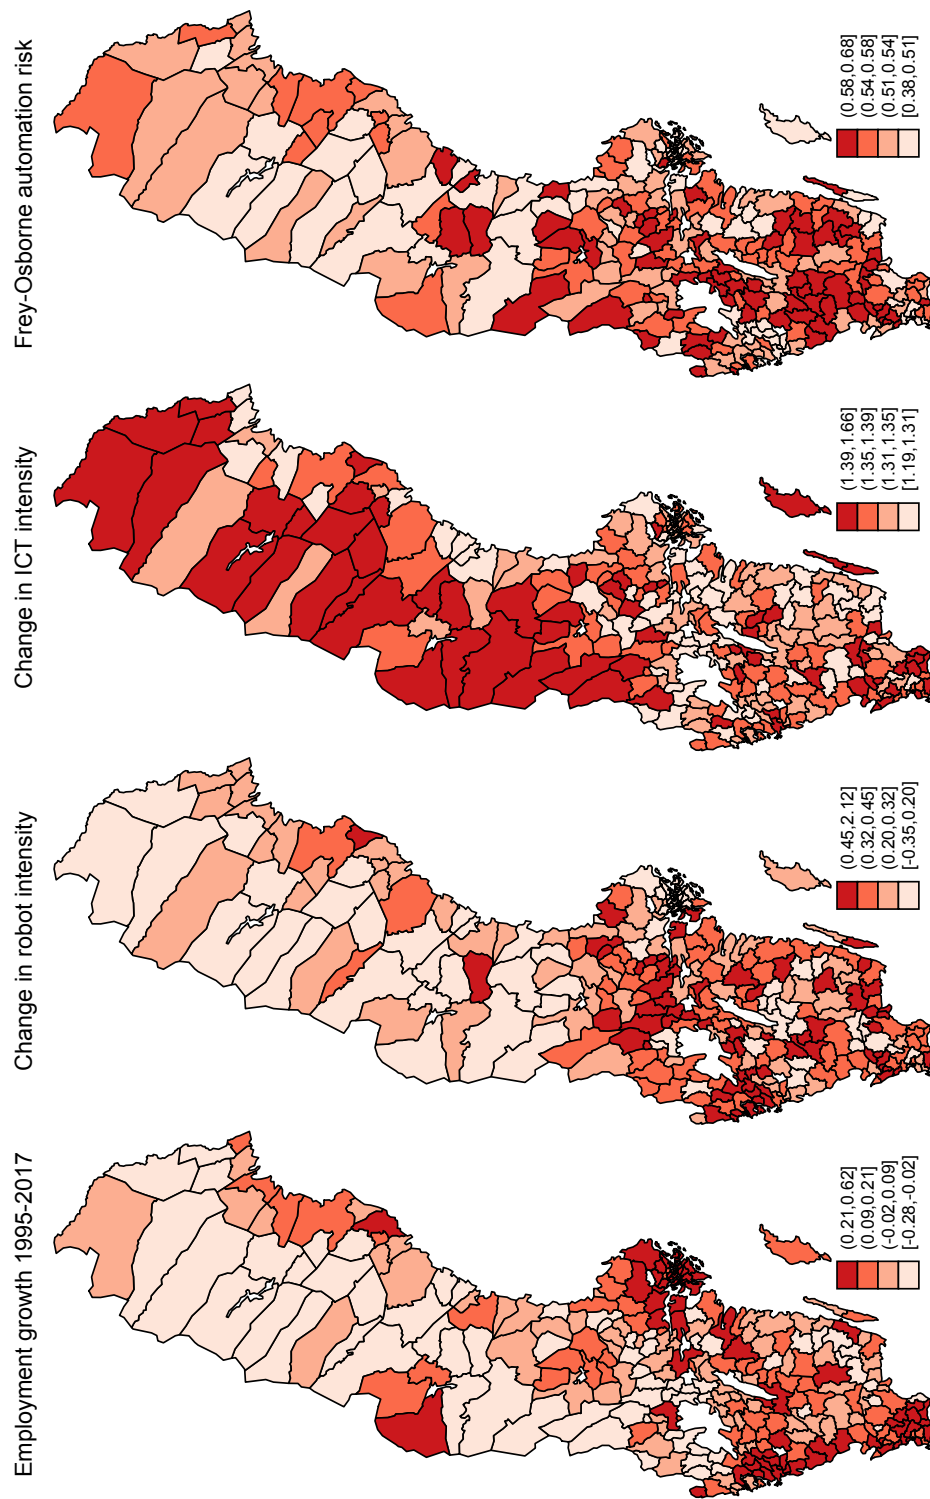
- EGGERTSSON, G. B., J. A. ROBBINS, AND E. G. WOLD (2018): “Kaldor and Pikettys Facts: The Rise of Monopoly Power in the United States,” Working Paper 24287, National Bureau of Economic Research.
- ELIASON, M., AND D. STORRIE (2006): “Lasting or Latent Scars? Swedish Evidence on the Long-Term Effects of Job Displacement,” *Journal of Labor Economics*, 24(4), 831–856.
- ELSBY, M., B. HOBIJN, AND A. SAHIN (2013): “The Decline of the U.S. Labor Share,” *Brookings Papers on Economic Activity*, 44(2 (Fall)), 1–63.
- FENG, A., AND G. GRAETZ (forthcoming): “Training Requirements, Automation, and Job Polarization,” *Economic Journal*.
- FLAM, H. (2019): “Det är facket, dumbom!,” blog post at *Ekonomistas*, 3 March.
- FORD, M. (2015): *Rise of the Robots: Technology and the Threat of a Jobless Future*. Basic Books.
- FREY, C. B. (2019): *The Technology Trap: Capital, Labor, and Power in the Age of Automation*. Princeton University Press.
- FREY, C. B., AND M. A. OSBORNE (2017): “The future of employment: How susceptible are jobs to computerisation?,” *Technological Forecasting and Social Change*, 114, 254 – 280.
- GAGGL, P., AND G. C. WRIGHT (2017): “A Short-Run View of What Computers Do: Evidence from a UK Tax Incentive,” *American Economic Journal: Applied Economics*, 9(3), 262–94.
- GOOS, M., AND A. MANNING (2007): “Lousy and Lovely Jobs: The Rising Polarization of Work in Britain,” *The Review of Economics and Statistics*, 89(1), 118–133.
- GOOS, M., A. MANNING, AND A. SALOMONS (2014): “Explaining Job Polarization: Routine-Biased Technological Change and Offshoring,” *American Economic Review*, 104(8), 2509–26.
- GORDON, R. J. (2012): “Is U.S. Economic Growth Over? Faltering Innovation Confronts the Six Headwinds,” Working Paper 18315, National Bureau of Economic Research.

- (2014): “The Demise of U.S. Economic Growth: Restatement, Rebuttal, and Reflections,” Working Paper 19895, National Bureau of Economic Research.
- GRAETZ, G. (2019): “Labor Demand in the Past, Present, and Future,” European Economy Discussion Papers 114, European Commission.
- (2020): “Technological Change and the Swedish Labor Market,” Report to the Swedish Fiscal Policy Council 2020/1.
- GRAETZ, G., AND G. MICHAELS (2018): “Robots at Work,” *The Review of Economics and Statistics*, 100(5), 753–768.
- GREGG, P., S. MACHIN, AND M. FERNANDEZ-SALGADO (2014): “Real Wages and Unemployment in the Big Squeeze,” *The Economic Journal*, 124(576), 408–432.
- GUTIERREZ, G., AND S. PITON (2019): “Revisiting the Global Decline of the (Non-Housing) Labor Share,” Discussion Papers 1913, Centre for Macroeconomics (CFM).
- HAGELUND, K., E. W. NORDBØ, AND L. SAUVIK (2017): “Lønnsandelen,” Aktuell kommentar, Norges Bank.
- HENSVIK, L., AND O. N. SKANS (2019): “The skill-specific impact of past and projected occupational decline,” Working Paper Series 2019:28, IFAU – Institute for Evaluation of Labour Market and Education Policy.
- HERRENDORF, B., R. ROGERSON, AND A. VALENTINYI (2014): “Chapter 6 - Growth and Structural Transformation,” in *Handbook of Economic Growth*, ed. by P. Aghion, and S. N. Durlauf, vol. 2 of *Handbook of Economic Growth*, pp. 855 – 941. Elsevier.
- HSIEH, C.-T., E. HURST, C. I. JONES, AND P. J. KLENOW (2019): “The Allocation of Talent and U.S. Economic Growth,” *Econometrica*, 87(5), 1439–1474.
- HUMLUM, A. (2019): “Robot Adoption and Labor Market Dynamics,” mimeo.
- INTERNATIONAL FEDERATION OF ROBOTICS (2019): “World Robotics,” report.

- JACOBSON, L. S., R. J. LALONDE, AND D. G. SULLIVAN (1993): “Earnings Losses of Displaced Workers,” *The American Economic Review*, 83(4), 685–709.
- JONES, C. (2016): “Chapter 1 - The Facts of Economic Growth,” vol. 2 of *Handbook of Macroeconomics*, pp. 3 – 69. Elsevier.
- KARABARBOUNIS, L., AND B. NEIMAN (2014): “The Global Decline of the Labor Share,” *The Quarterly Journal of Economics*, 129(1), 61–103.
- (2018): “Accounting for Factorless Income,” NBER Working Papers 24404, National Bureau of Economic Research, Inc.
- KATZ, L. F., AND K. M. MURPHY (1992): “Changes in Relative Wages, 1963–1987: Supply and Demand Factors,” *The Quarterly Journal of Economics*, 107(1), 35–78.
- KOCH, M., I. MANUYLOV, AND M. SMOLKA (2019): “Robots and Firms,” Economics Working Papers 2019-05, Department of Economics and Business Economics, Aarhus University.
- KONJUNKTURINSTITUTET (2018): “Lönebildningsrapporten,” report.
- KRUEGER, A. B. (1993): “How Computers Have Changed the Wage Structure: Evidence from Microdata, 1984–1989,” *The Quarterly Journal of Economics*, 108(1), 33–60.
- KRUSELL, P., L. E. OHANIAN, J.-V. ROS-RULL, AND G. L. VIOLANTE (2000): “Capital-Skill Complementarity and Inequality: A Macroeconomic Analysis,” *Econometrica*, 68(5), 1029–1054.
- LINDLEY, J., AND S. MACHIN (2016): “The Rising Postgraduate Wage Premium,” *Economica*, 83(330), 281–306.
- LIPSIUS, B. (2018): “Labor Market Concentration does not Explain the Falling Labor Share,” 2018 Papers pli1202, Job Market Papers.
- MACHIN, S., AND J. VAN REENEN (1998): “Technology and Changes in Skill Structure: Evidence from Seven OECD Countries,” *The Quarterly Journal of Economics*, 113(4), 1215–1244.

- MANN, K., AND L. PÜTTMANN (2017): mimeo.
- MICHAELS, G., A. NATRAJ, AND J. VAN REENEN (2014): “Has ICT Polarized Skill Demand? Evidence from Eleven Countries over Twenty-Five Years,” *The Review of Economics and Statistics*, 96(1), 60–77.
- MICHAELS, G., F. RAUCH, AND S. J. REDDING (2018): “Task Specialization in U.S. Cities from 1880 to 2000,” *Journal of the European Economic Association*, 17(3), 754–798.
- MITCHELL, M. (2019): *Artificial Intelligence: A Guide for Thinking Humans*. Farrar, Straus and Giroux.
- MORAVEC, H. (1988): *Mind Children*. Harvard University Press.
- MORETTI, E. (2010): “Local Multipliers,” *American Economic Review*, 100(2), 373–77.
- NORDHAUS, W. D. (2007): “Two Centuries of Productivity Growth in Computing,” *The Journal of Economic History*, 67(01), 128–159.
- OECD (2019): *OECD Employment Outlook 2019: The Future of Work*. OECD Publishing, Paris.
- OECD (2020): “Educational attainment and labour-force status,” database.
- O’MAHONY, M., AND M. P. TIMMER (2009): “Output, Input and Productivity Measures at the Industry Level: The EU KLEMS Database,” *Economic Journal*, 119(538), F374–F403.
- POLANYI, M. (1966): *The Tacit Dimension*. New York: Doubleday.
- RINZ, K. (2018): “Labor Market Concentration, Earnings Inequality, and Earnings Mobility,” CARRA Working Papers 2018-10, Center for Economic Studies, U.S. Census Bureau.
- ROSSI-HANSBERG, E., P.-D. SARTE, AND N. TRACHTER (2018): “Diverging Trends in National and Local Concentration,” Working Paper 25066, National Bureau of Economic Research.

- SEMUELS, A. (2017): “Why Does Sweden Have So Many Start-Ups?,” *The Atlantic*.
- SHILLER, R. J. (2019): “Narratives about technology-induced job degradation then and now,” *Journal of Policy Modeling*.
- SOMERS, J. (2017): “Is AI Riding a One-Trick Pony?,” *MIT Technology Review*.
- TRAINA, J. (2018): “Is Aggregate Market Power Increasing? Production Trends Using Financial Statements,” mimeo.
- VENERI, C. M. (1997): “Evaluating the 1995 occupational employment projections,” *Monthly Labor Review*, 120.
- VERDUGO, G., H. FRAISSE, AND G. HORNY (2012): “Changes in Wage Inequality in France,” *Revue économique*, 63(6), 1081–1112.
- WEBB, M. (2020): “The Impact of Artificial Intelligence on the Labor Market,” mimeo.
- WRIGHT, E. O., AND R. E. DWYER (2003): “The patterns of job expansions in the USA: a comparison of the 1960s and 1990s,” *Socio-Economic Review*, 1(3), 289–325.



Notes: Employment growth is measured in logarithmic units (weighted mean across regions: 0.23). Robot intensity is defined as robots per million hours worked (weighted mean change: 0.31). The change in ICT intensity is the log change in ICT capital minus the log change in hours worked (weighted mean: 1.35). Frey-Osborne automation risk is the probability that a worker's occupation becomes automatable (weighted mean: 0.53).

Figure 10: Employment growth and technological change across Swedish municipalities

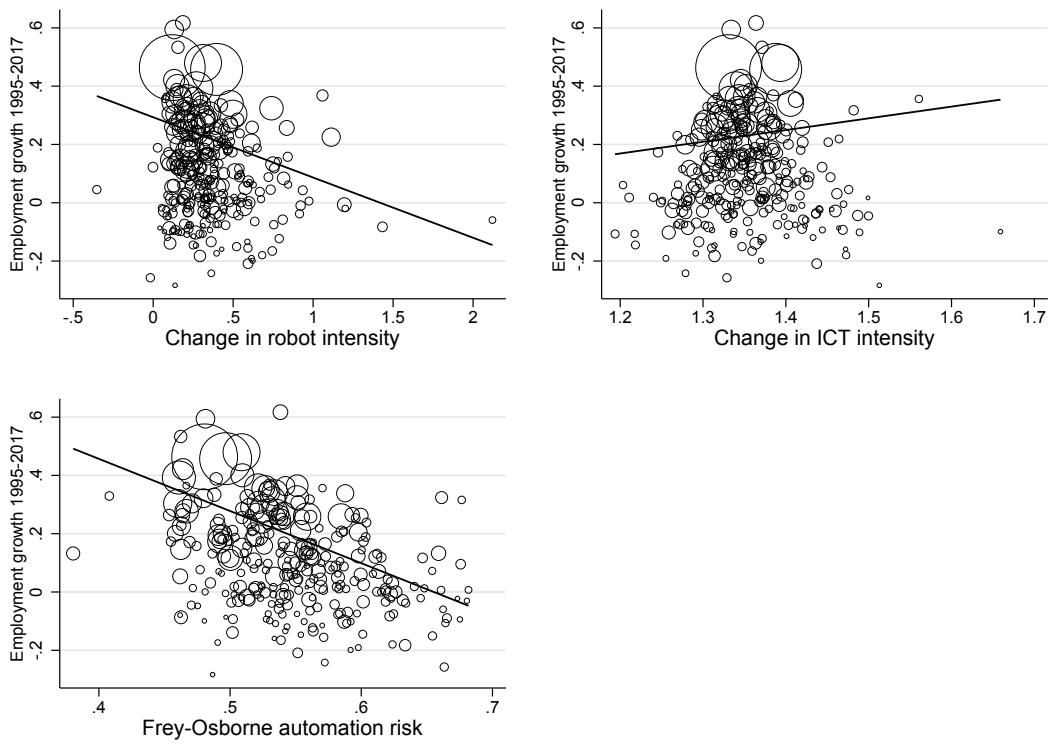
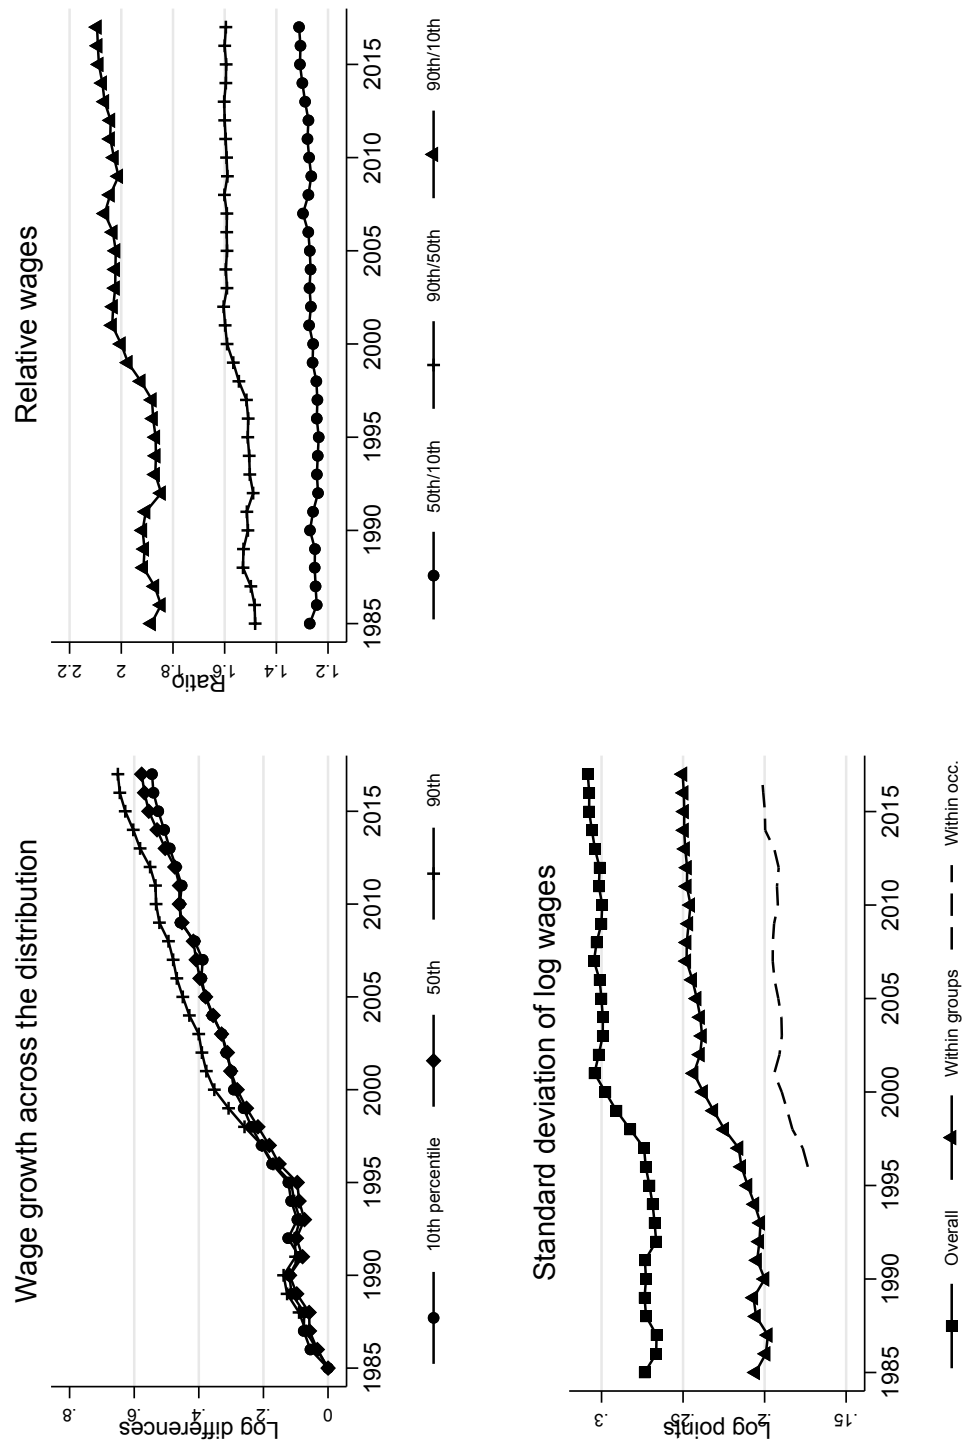


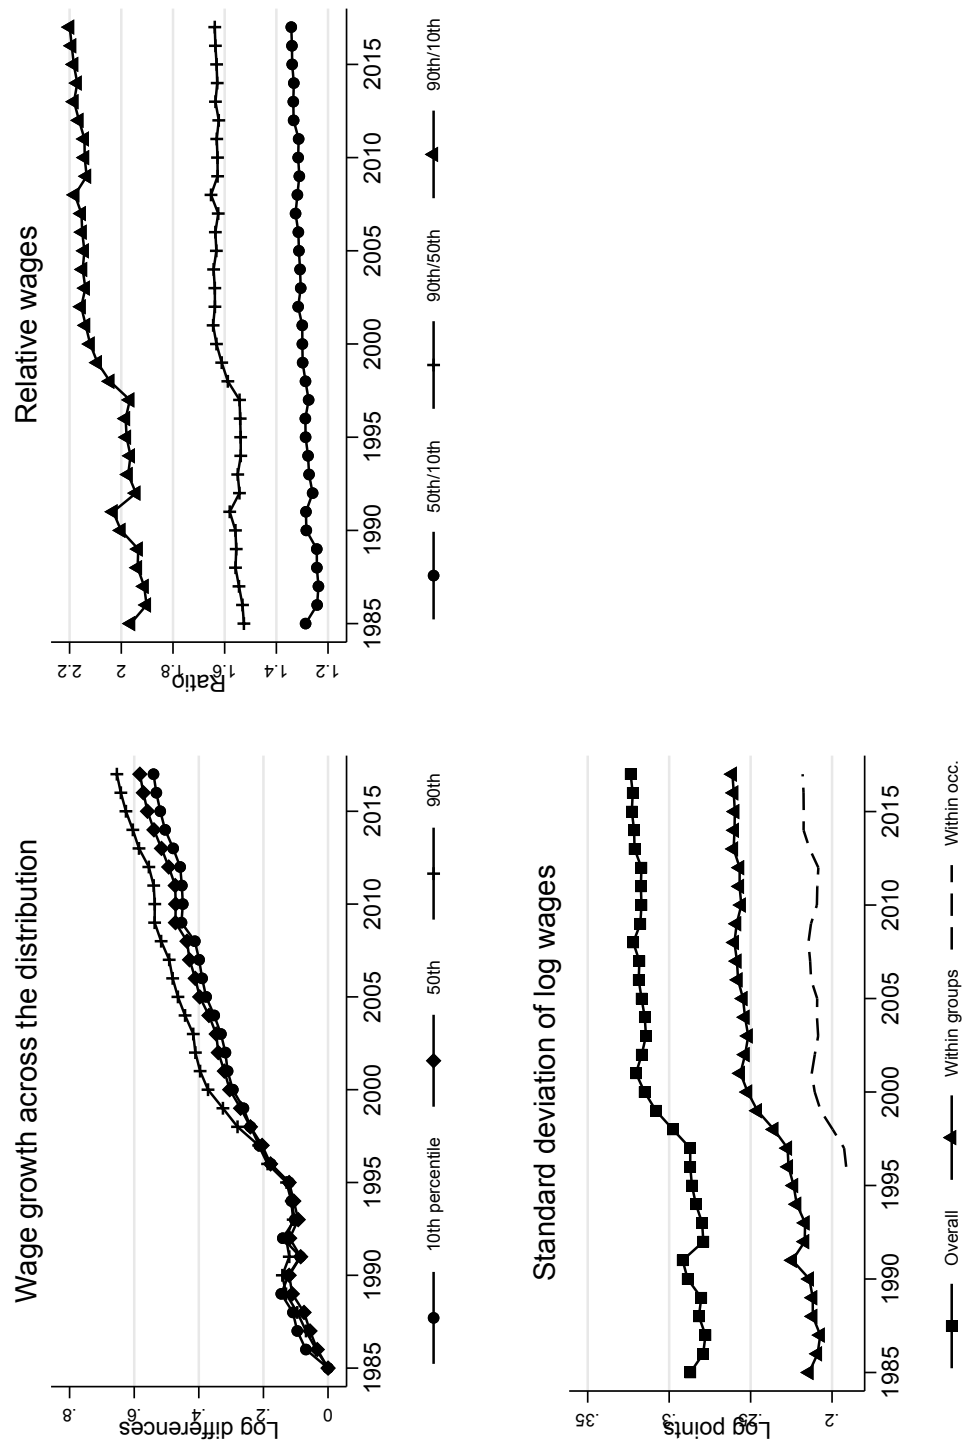
Figure 11: Employment growth and technological change across Swedish municipalities

Appendix figures and tables



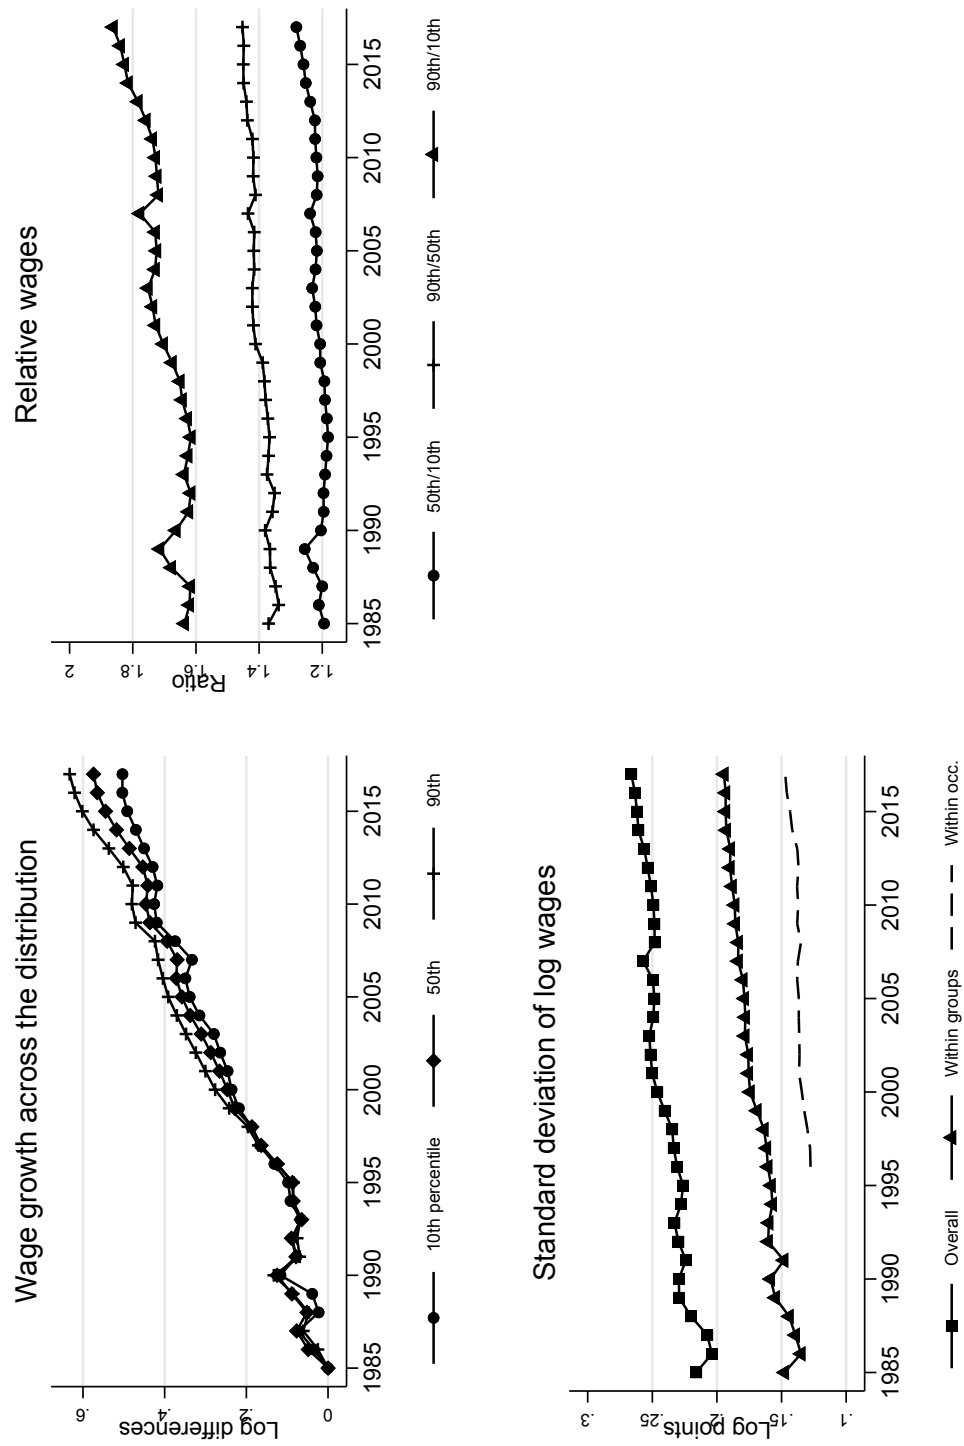
Sample: Full-time employees aged 16-64

Figure A1: Wage growth and wage inequality—full-time workers



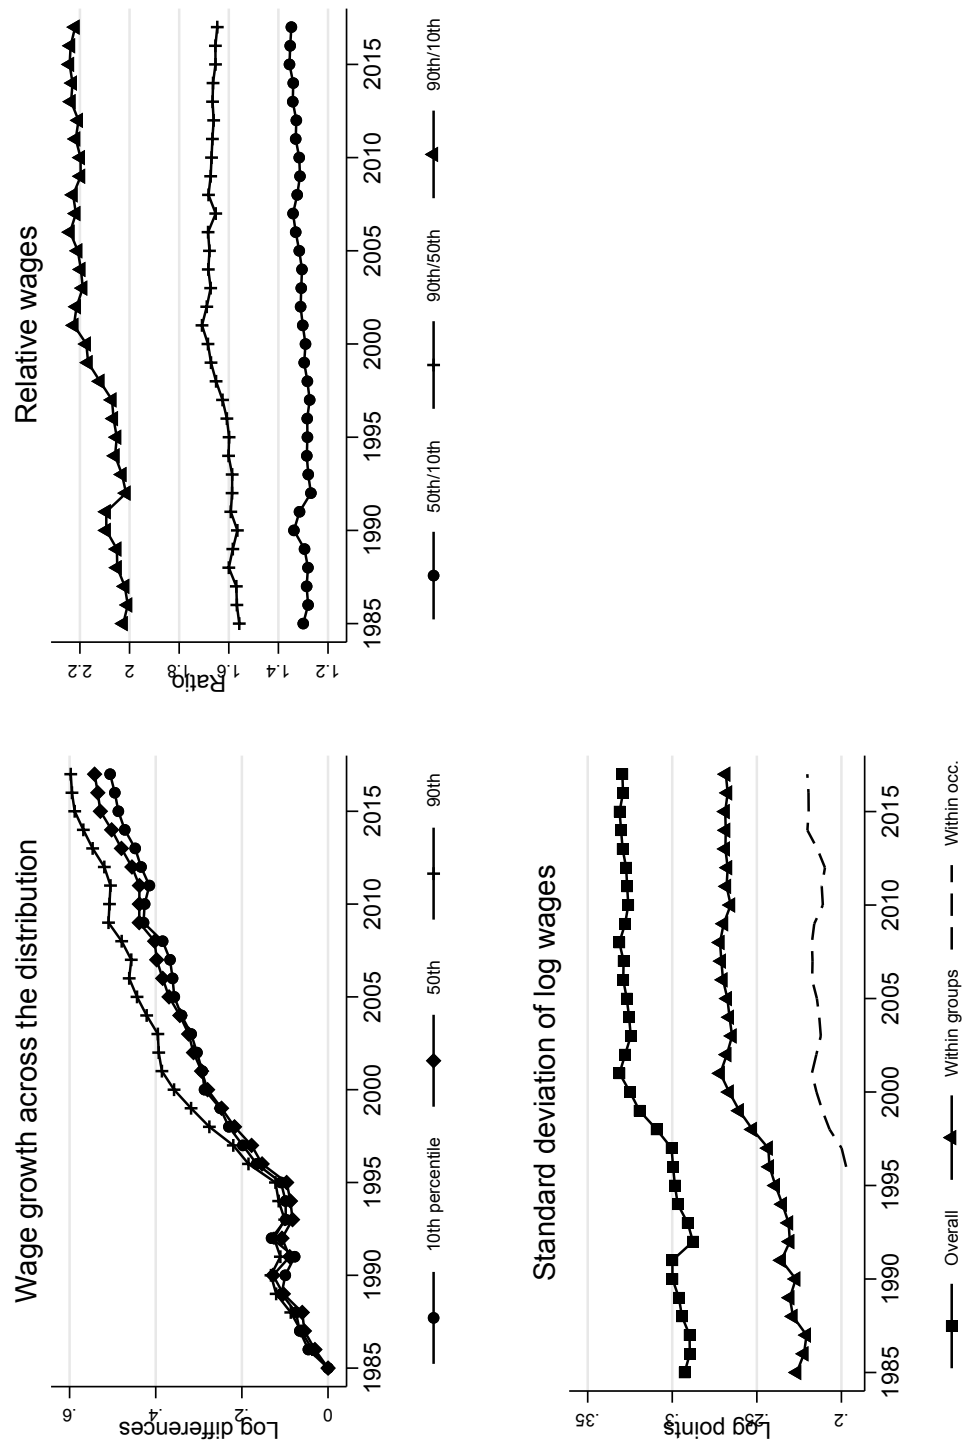
Sample: Private sector employees aged 16-64

Figure A2: Wage growth and wage inequality—private sector



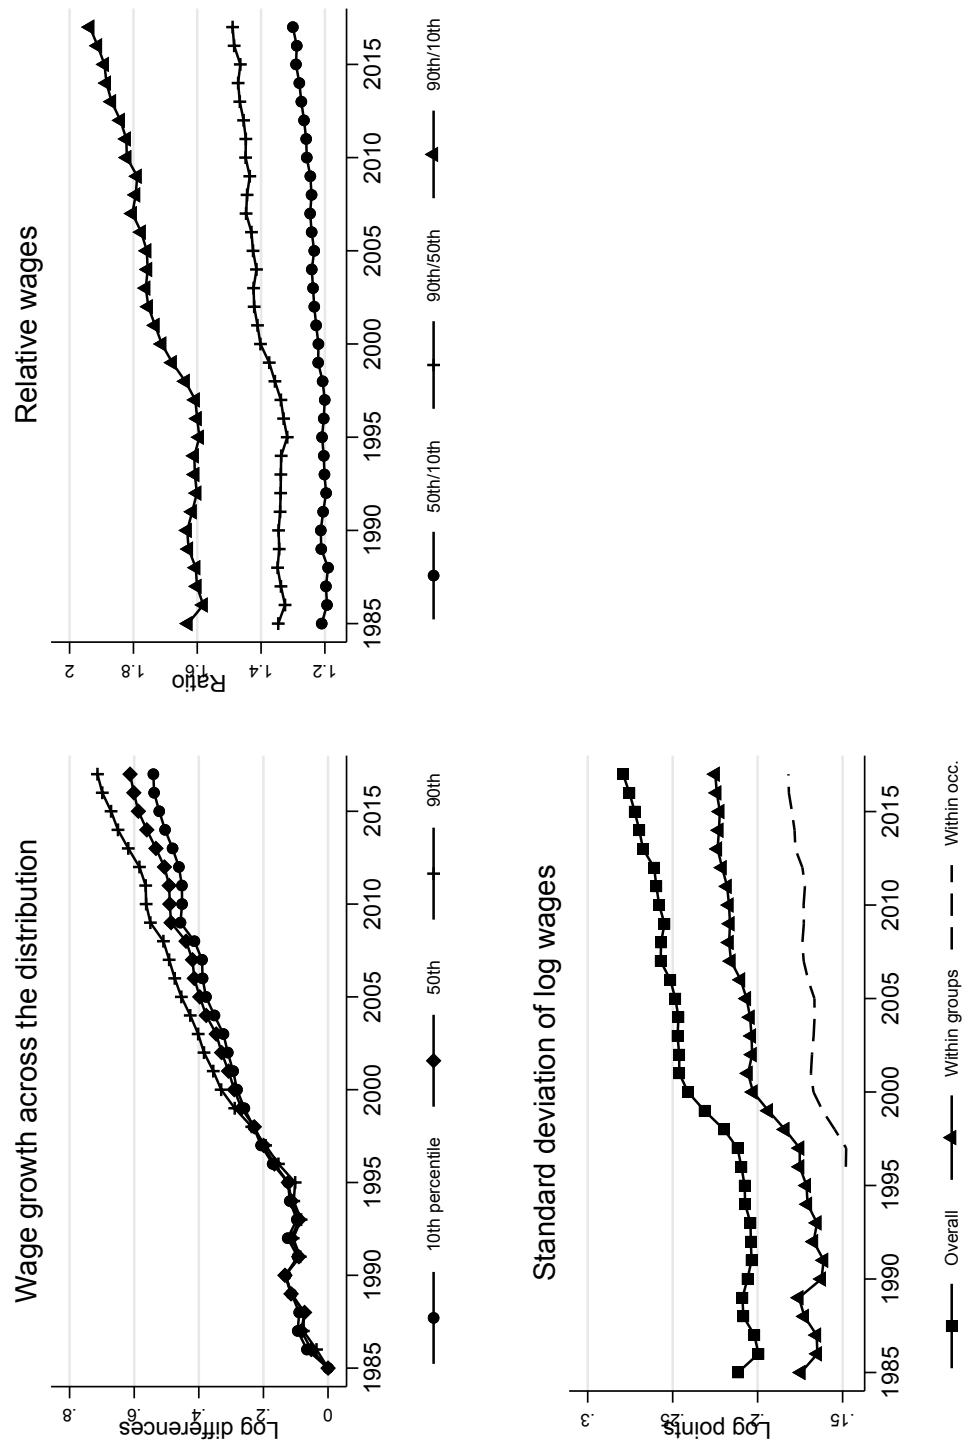
Sample: Public sector employees aged 16-64

Figure A3: Wage growth and wage inequality—public sector



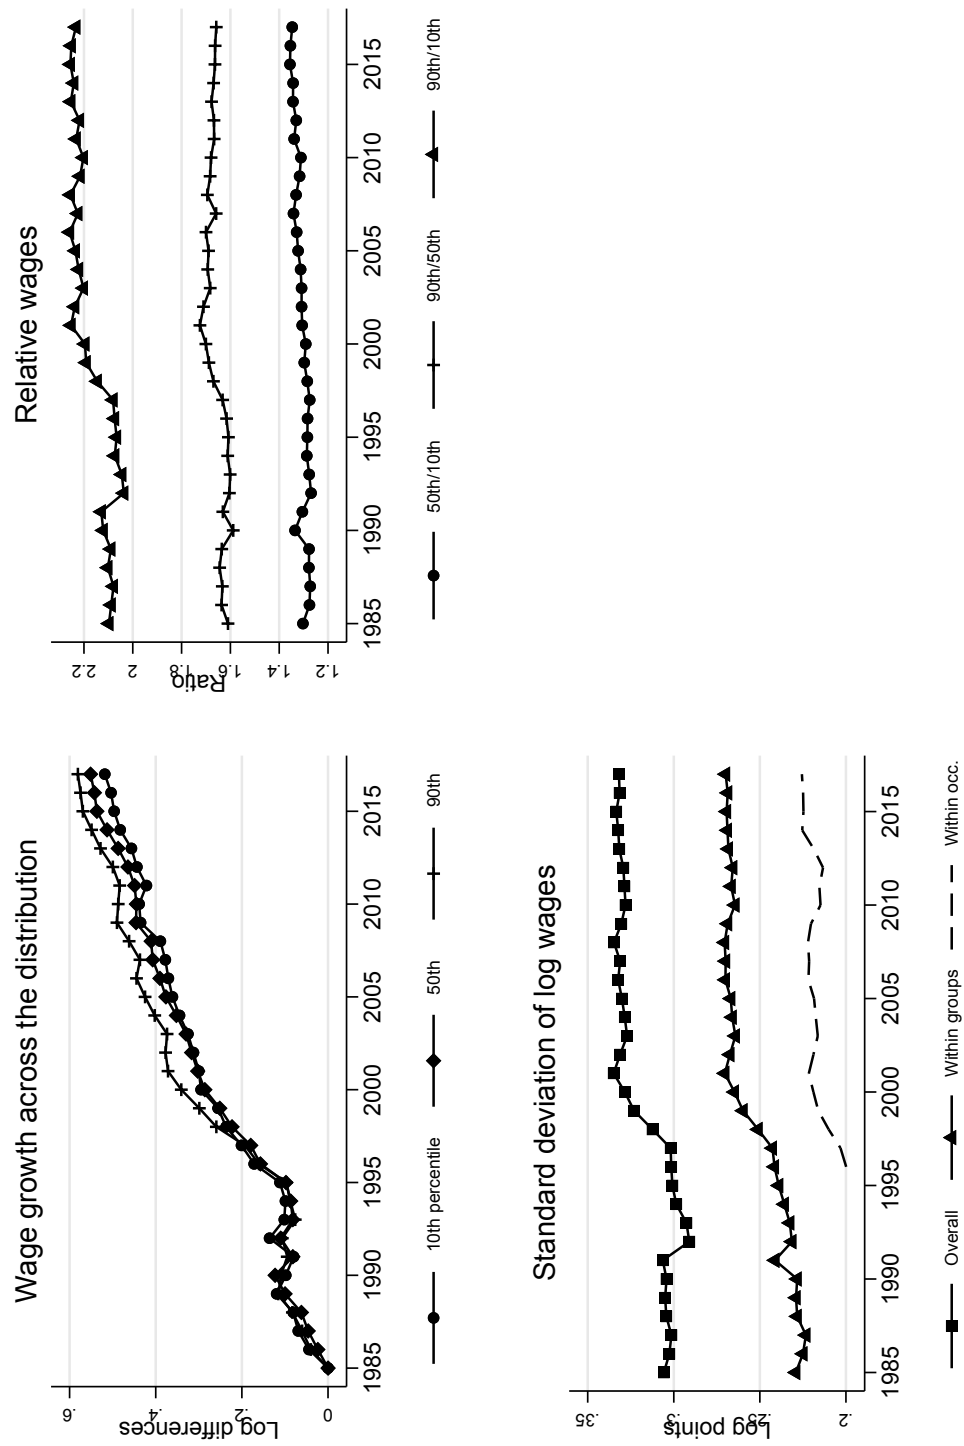
Sample: Men aged 16-64

Figure A4: Wage growth and wage inequality—men



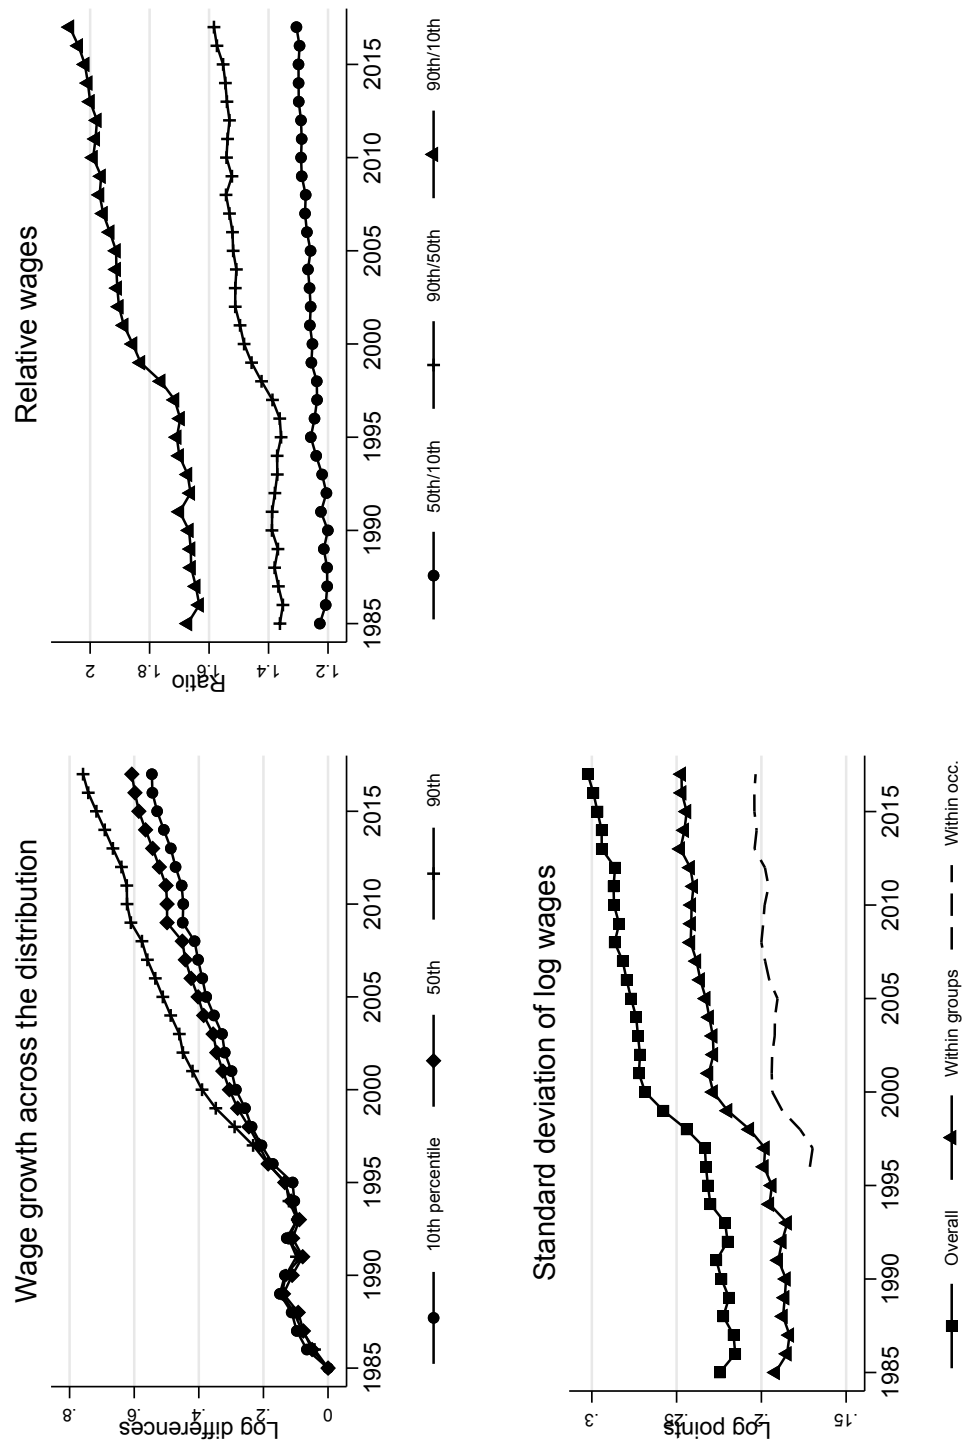
Sample: Women aged 16-64

Figure A5: Wage growth and wage inequality—women



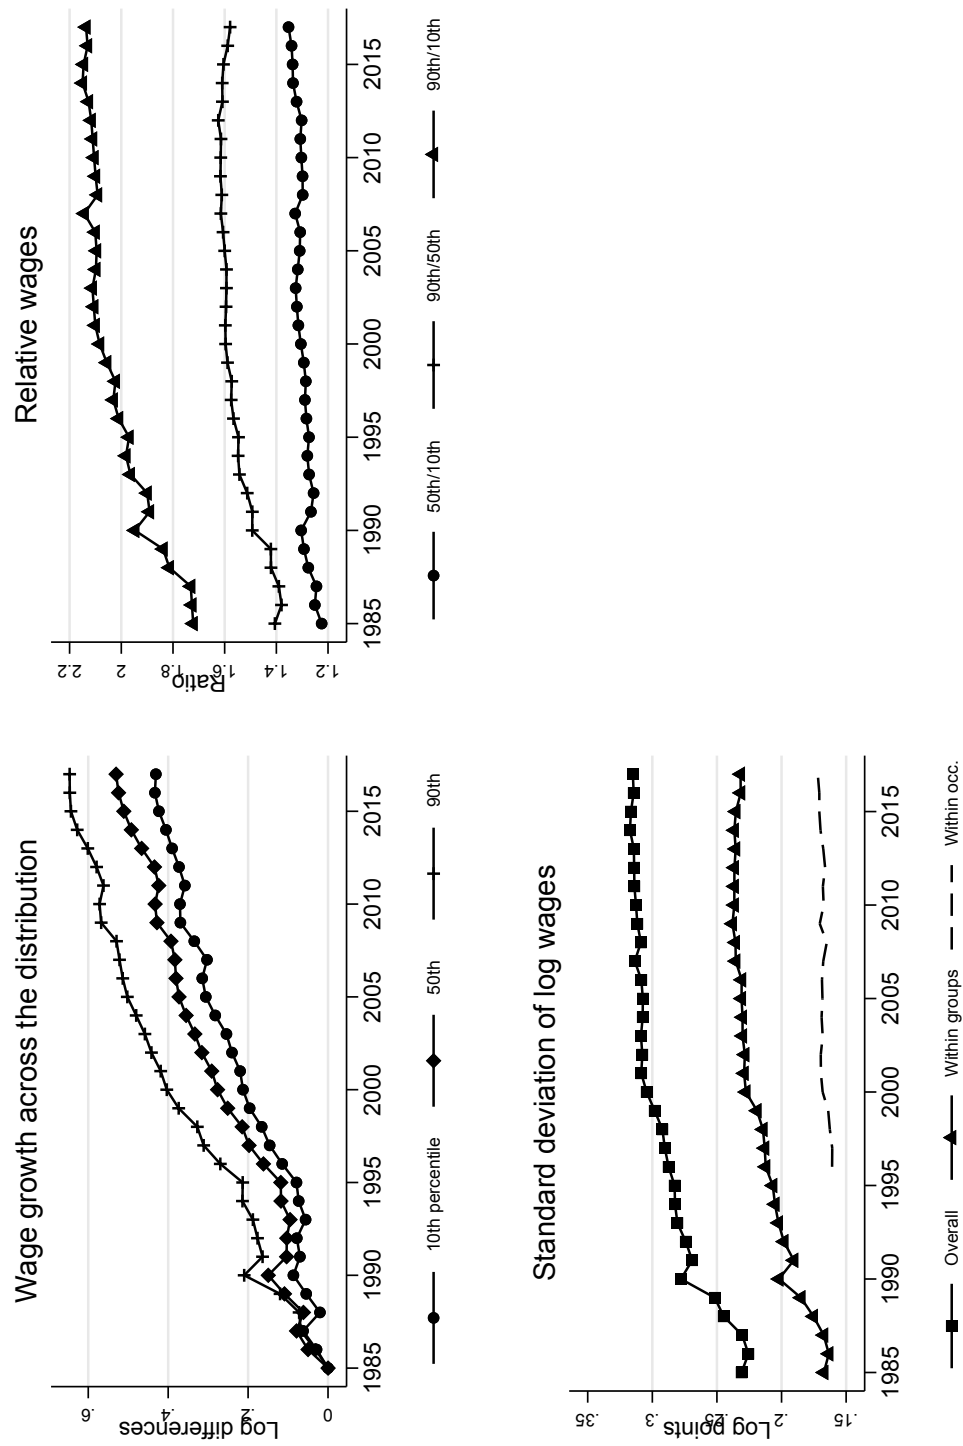
Sample: Male private sector employees aged 16-64

Figure A6: Wage growth and wage inequality—private sector, men



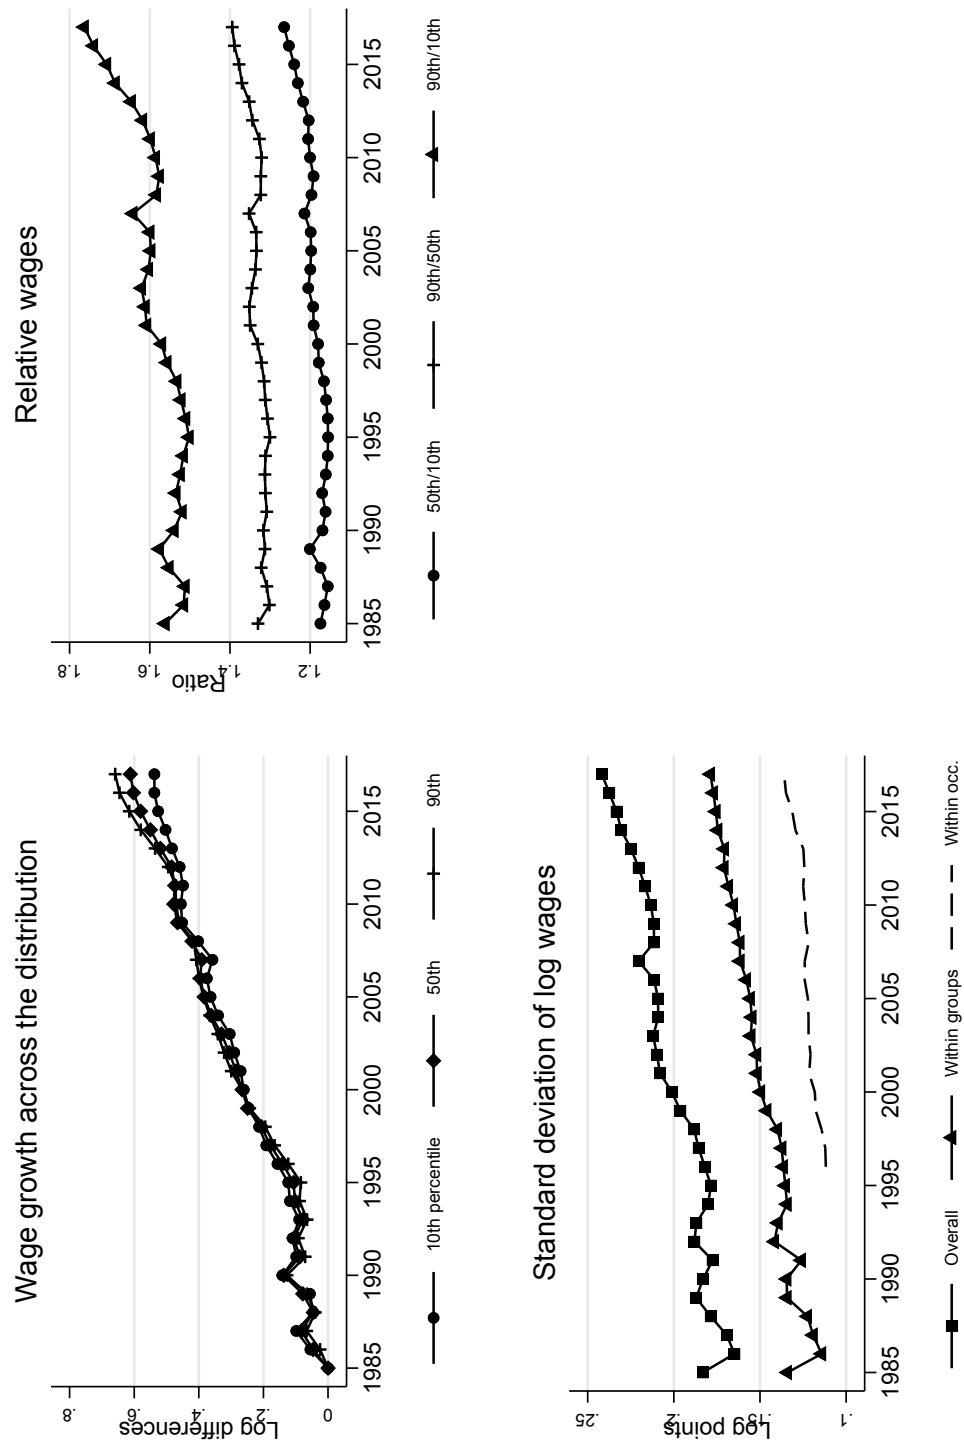
Sample: Female private sector employees aged 16-64

Figure A7: Wage growth and wage inequality—private sector, women



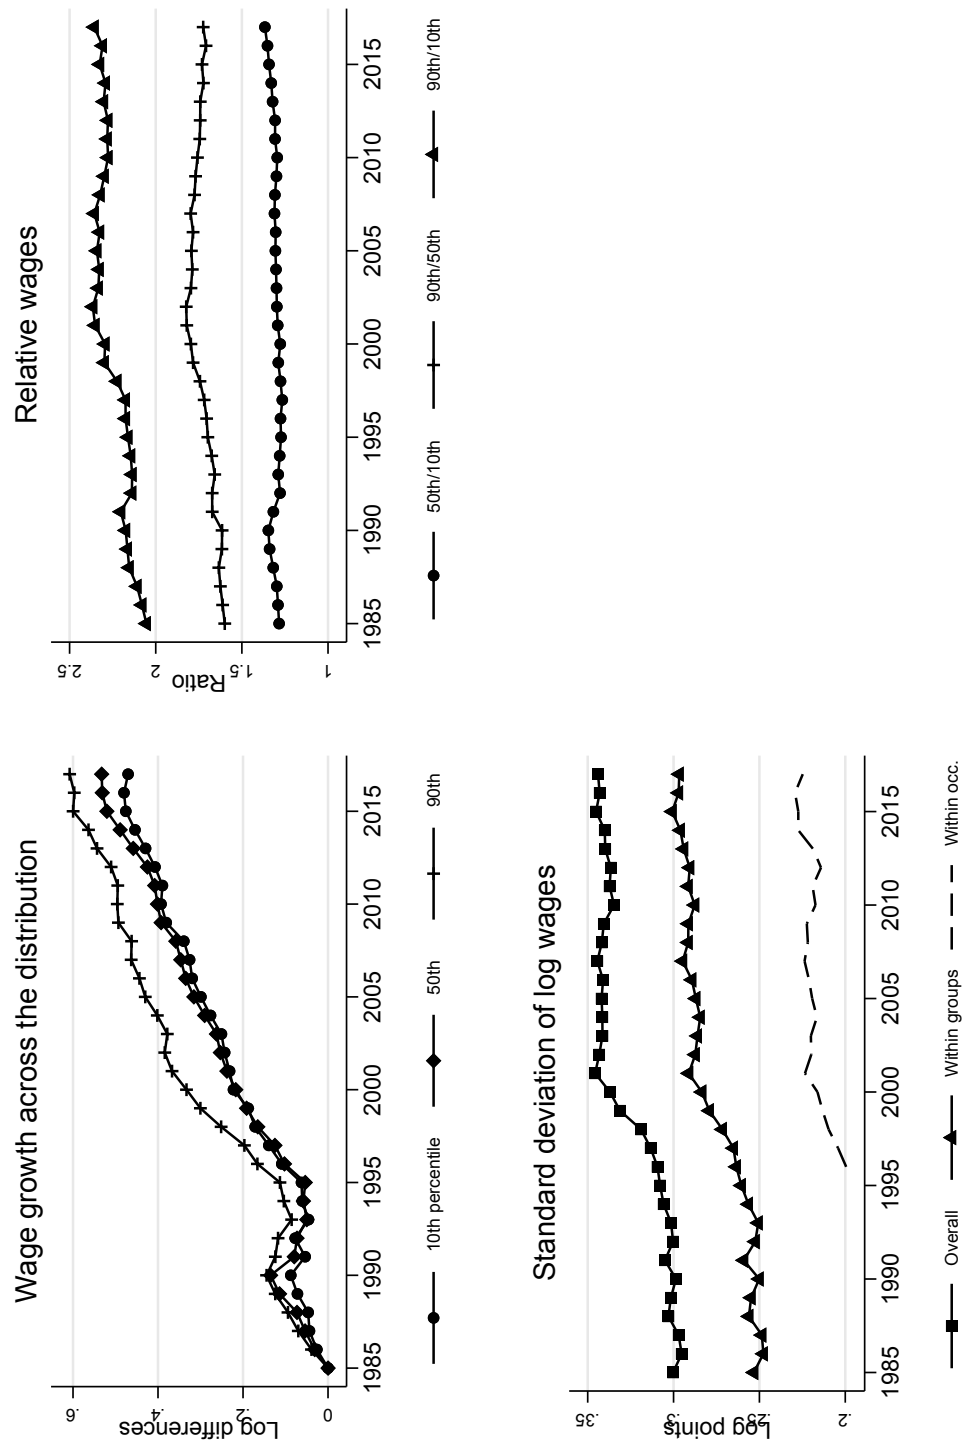
Sample: Male public sector employees aged 16-64

Figure A8: Wage growth and wage inequality—public sector, men



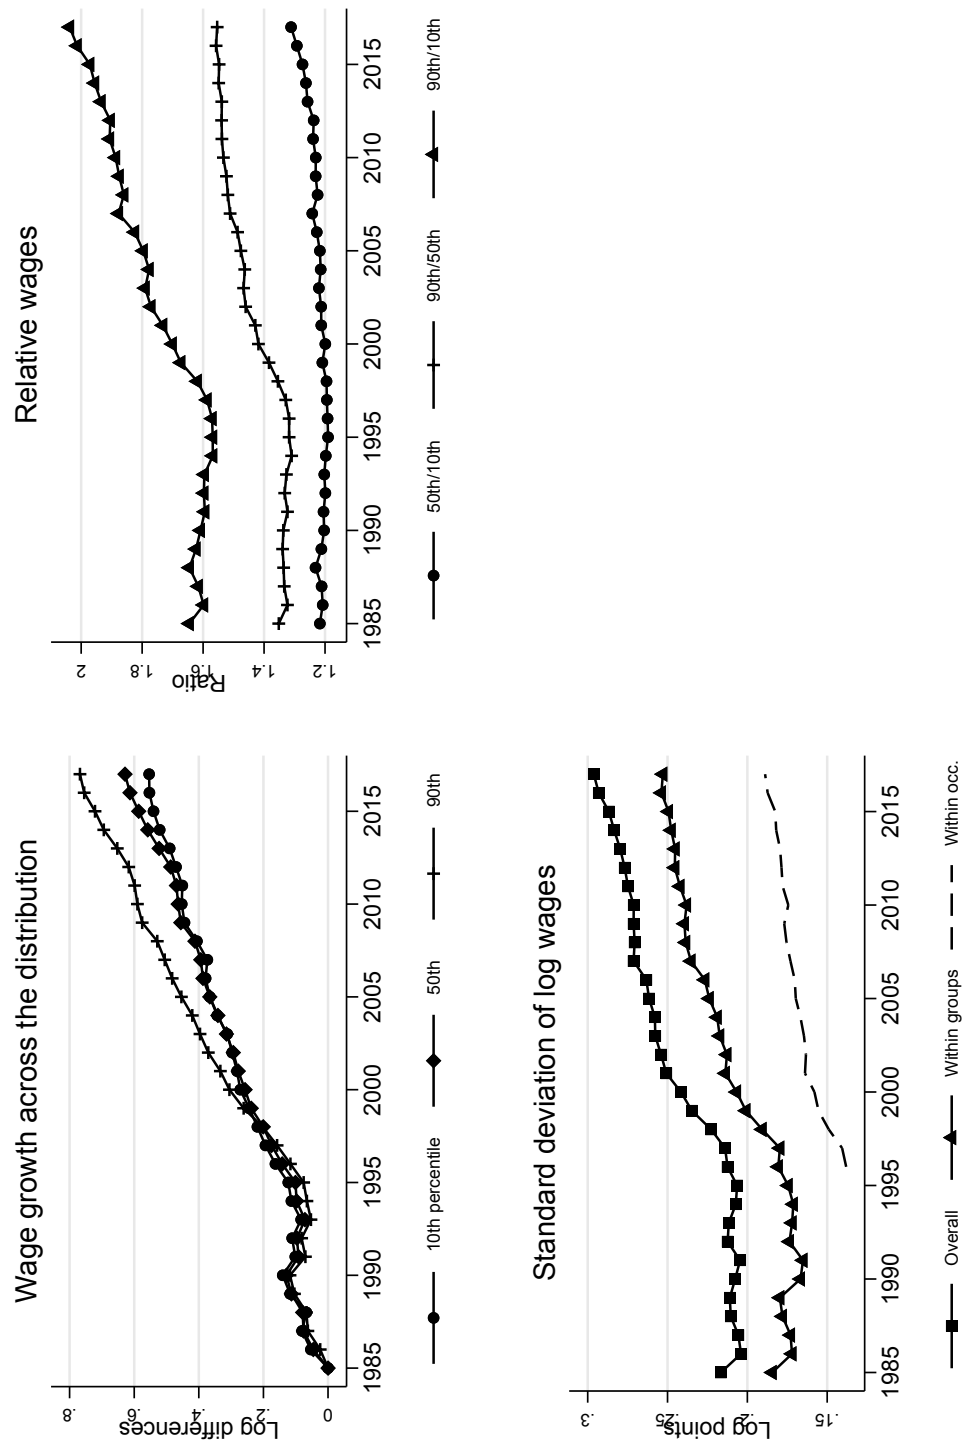
Sample: Female public sector employees aged 16-64

Figure A9: Wage growth and wage inequality—public sector, women



Sample: Men aged 40-49, employed full time

Figure A10: Wage growth and wage inequality—prime-age men



Sample: Women aged 40-49, employed full time

Figure A11: Wage growth and wage inequality—prime-age women

Table A1: Overview of occupational classification

2-digit category	ISCO
Officials & Managers	Legislators & senior officials Corporate managers Managers of small enterprises
Professionals	Physical, mathematical & engineering science professionals Life science & health professionals Teaching professionals Other professionals
Technicians	Physical & engineering science associate professionals Life science & health associate professionals Teaching associate professionals Other associate professionals Skilled agricultural & fishery workers
Clerks	Office clerks Customer services clerks
Service & Sales	Personal & protective services workers Models, salespersons & demonstrators
Elementary Occupations	Sales & services elementary occupations Agricultural, fishery & related labourers Labourers in mining, construction, manufacturing & transport
Crafts	Extraction & building trades workers Metal, machinery & related trades workers Precision, handicraft, craft printing & related trades workers Other craft & related trades workers
Operators & Assemblers	Stationary-plant & related operators Machine operators & assemblers Drivers & mobile plant operators

Table A2: Technological change across regions: Skill share changes (wage bill)

	(1)	(2)	(3)
<i>A. Change in high-skill wage bill share 1995-2017 (weighted mean: 0.13)</i>			
Robots	-0.00092 (0.0100)	-0.00039 (0.0100)	-0.019 (0.012)
ICT	0.11 (0.052)	0.12 (0.049)	0.16 (0.049)
<i>A. Change in middle-skill wage bill share 1995-2017 (weighted mean: -0.02)</i>			
Robots	0.026 (0.011)	-0.00094 (0.0081)	-0.0098 (0.011)
ICT	-0.10 (0.057)	-0.13 (0.046)	-0.11 (0.046)
<i>A. Change in low-skill wage bill share 1995-2017 (weighted mean: -0.11)</i>			
Robots	-0.024 (0.0086)	0.0030 (0.0069)	0.029 (0.0090)
ICT	-0.018 (0.039)	-0.00042 (0.018)	-0.054 (0.018)
Other capital	✓	✓	✓
Employment and population controls	✓	✓	✓
Education controls		✓	✓
Industry controls			✓

Notes: Results are from regressions of the dependent variables indicated in the panel headings on robot adoption, the change in ICT intensity, and controls. The sample consists of 284 Swedish municipalities. ‘Other capital’ indicates that the change in non-ICT intensity is controlled for. Employment and population controls include the log of the employment-to-population ratio as well as the log of employment in 1990. Education controls include the employment shares of three educational groups, and industry controls include the employment shares of four broad sectoral groups, both in 1990. Regressions are weighted by initial employment shares. Robust standard errors in parentheses.

Table A3: Technological change across regions: Changes in sectoral employment shares

	(1)	(2)	(3)
<i>A. Change in manufacturing employment share (weighted mean: -0.08)</i>			
Robots	-0.060 (0.016)	-0.048 (0.017)	-0.00066 (0.018)
ICT	0.16 (0.053)	0.17 (0.051)	0.082 (0.051)
<i>B. Change in services employment share (weighted mean: 0.09)</i>			
Robots	0.049 (0.016)	0.040 (0.018)	-0.0017 (0.016)
ICT	-0.093 (0.053)	-0.095 (0.052)	-0.043 (0.051)
<i>C. Change in utilities employment share (weighted mean: -0.02)</i>			
Robots	0.0048 (0.0038)	0.0049 (0.0043)	-0.0059 (0.0047)
ICT	-0.071 (0.021)	-0.069 (0.020)	-0.016 (0.014)
<i>D. Change in primary employment share (weighted mean: 0.01)</i>			
Robots	-0.00026 (0.0038)	-0.0017 (0.0041)	0.014 (0.0047)
ICT	-0.0040 (0.022)	-0.010 (0.020)	-0.045 (0.020)
Other capital	✓	✓	✓
Employment and population controls	✓	✓	✓
Education controls		✓	✓
Industry controls			✓

Notes: Results are from regressions of the dependent variables indicated in the panel headings on robot adoption, the change in ICT intensity, and controls. The sample consists of 284 Swedish municipalities. ‘Other capital’ indicates that the change in non-ICT intensity is controlled for. Employment and population controls include the log of the employment-to-population ratio as well as the log of employment in 1990. Education controls include the employment shares of three educational groups, and industry controls include the employment shares of four broad sectoral groups, both in 1990. Regressions are weighted by initial employment shares. Robust standard errors in parentheses.

Table A4: Technological change across regions: Recent employment growth and automation risk

	(1)	(2)	(3)	(4)
<i>Employment growth 1995-2017 (weighted mean: 0.23)</i>				
Frey-Osborne automation risk	-1.79 (0.34)	-0.53 (0.16)	-0.37 (0.22)	-0.24 (0.21)
Employment and population controls		✓	✓	✓
Education controls			✓	✓
Industry controls				✓

Notes: Results are from regressions of employment growth 1995-2017 on the Frey and Osborne (2017) measure of automation risk and controls. The sample consists of 284 Swedish municipalities. Employment and population controls include the log of the employment-to-population ratio as well as the log of employment in 1990. Education controls include the employment shares of three educational groups, and industry controls include the employment shares of four broad sectoral groups, both in 1990. Regressions are weighted by initial employment shares. Robust standard errors in parentheses.