University of Groningen

Thriving in Times of Technological Change

Cnossen, Femke

DOI:
[10.33612/diss.204617165](https://doi.org/10.33612/diss.204617165)

IMPORTANT NOTE: You are advised to consult the publisher's version (publisher's PDF) if you wish to cite from it. Please check the document version below.

Document Version
Publisher's PDF, also known as Version of record

Publication date:
2022

[Link to publication in University of Groningen/UMCG research database](#)

Citation for published version (APA):
Cnossen, F. (2022). *Thriving in Times of Technological Change: how tasks, skills and meaning shape careers in the 21st century labour market*. [Thesis fully internal (DIV), University of Groningen]. University of Groningen, SOM research school. <https://doi.org/10.33612/diss.204617165>

Copyright

Other than for strictly personal use, it is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license (like Creative Commons).

The publication may also be distributed here under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license. More information can be found on the University of Groningen website: <https://www.rug.nl/library/open-access/self-archiving-pure/taverne-amendment>.

Take-down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Downloaded from the University of Groningen/UMCG research database (Pure): <http://www.rug.nl/research/portal>. For technical reasons the number of authors shown on this cover page is limited to 10 maximum.

Thriving in Times of Technological Change

How tasks, skills and meaning shape careers
in the 21st century labour market

Femke Cnossen

Cover design: Kiki Chang Szu Yun & Femke Cnossen
Cover image: Adobe Stock
Printing: Ridderprint B.V., www.ridderprint.nl

© Femke Cnossen, Groningen 2021. All rights reserved. No part of this thesis may be reproduced or transmitted in any form or by any means without prior written permission from the author. The copyright of the paper that has been published has been transferred to the respective journal.

rijksuniversiteit
 groningen

Thriving in Times of Technological Change

How tasks, skills and meaning shape careers
in the 21st century labour market

Proefschrift

ter verkrijging van de graad van doctor aan de
Rijksuniversiteit Groningen
op gezag van de
rector magnificus prof. dr. C. Wijmenga
en volgens besluit van het College voor Promoties.

De openbare verdediging zal plaatsvinden op

donderdag 10 maart 2022 om 14.30 uur

door

Femke Dina Cnossen

geboren op 19 september 1993
te Leeuwarden

Promotores

Prof. dr. R.C. Inklaar

Prof. dr. S. Brakman

Copromotor

Dr. M.V. Nikolova

Beoordelingscommissie

Prof. dr. R.J.M. Alessie

Prof. dr. M. Goos

Prof. dr. W. Dauth

Contents

Chapter 1	Introduction	7
Chapter 2	Technology, Tasks and Careers	29
Chapter 3	Learning the Right Skill	91
Chapter 4	Meaningful work	141
	Bibliography	183
	English summary	201
	Nederlandse samenvatting	209
	Acknowledgments/dankwoord	217
	About the author	221

Tasks
Skills
↓
Meaning

1

Introduction

“There’s never been a better time to be a worker with special skills or the right education, because these people can use technology to create and capture value. However, there’s never been a worse time to be a worker with only ‘ordinary’ skills and abilities to offer, because computers, robots, and other digital technologies are acquiring these skills and abilities at an extraordinary rate” (Brynjolfsson and McAfee, 2014, p.11).

The past decades can be characterised by a significant rise in the use of technology for the production of goods and services. By now, many of these technologies have obtained a secure foothold in the workplace: it is impossible to imagine an office without computers, healthcare without medical technology and factories without robots. This has heavily affected the organisation of work. Automation technologies have shown to replace programmeable and rules-oriented activities: routine tasks. However, when workers know how to use technology, it can complement their tasks and make them more productive. The replacing and augmenting nature of technology has led to growing inequalities between workers. Roughly speaking, those who have seen their tasks replaced are doing poorer, whereas those that have seen their tasks complemented are doing better.

This dissertation studies the changing nature of work. It centers around the main question: who thrives in labour markets that are increasingly shaped by technology? The thesis is split up in three chapters, each related to a different element of work and each taking a worker-level perspective: i) the tasks we do at work, ii) the skills we need to perform these tasks, and iii) the meaning we derive from our work.

Automation technologies can replace certain *tasks*, but they very rarely replace entire occupations. Nevertheless, some occupations contain relatively many replaceable – routine – tasks, which puts workers in such occupations at risk. Workers may adapt to new circumstances by taking up new tasks, and especially tasks that machines are not very good at. For instance, whereas accountants used to be predominantly focused on manual bookkeeping, the ‘modern’ accountant is heavily involved in face-to-face consultations with clients. This reduces the relative routine-intensity of the accountant’s tasks, and increases its nonroutine-intensity. In the first chapter of this dissertation, I study whether the relative (non)routine-intensity of one’s job correlates with their labour market outcomes: do people who execute relatively many replaceable tasks fare poorer on the labour market? And do those with nonroutine tasks perform better?

As technology changes the demand for certain tasks, so does the demand for *skills*: people use their abilities to execute tasks. Referring to the quote by Brynjolfsson and McAfee above: there has never been a better time to have special skills, and never a worse time to have ordinary skills. Their quote elicits a dichotomy of work: those who can work with the machine, and those who have to compete against it. So far, the literature has concluded that mainly the middle- and lower-educated workers have seen their tasks replaced, whereas higher-educated workers have been able to benefit from technology. This puts pressure on middle-educated young people that are currently entering the labour market: their jobs are disappearing. What should

this group of students learn in school, in order to ensure a successful start on today's labour market? In the second chapter of this dissertation, I analyse the curricula of all degrees in the Dutch vocational education system, to see whether we can find answers there: do students that learn certain skills at school do better than others?

Lastly, as more work can be automated, there is not only fear that our jobs can be replaced, but also that the work that is 'left for humans' is no longer *meaningful*. As meaningful work is not yet a prominent theme in the labour market literature, the last chapter argues that economists should care (more) about it. Especially in light of technological change, this theme is very likely to become more important in the future.

The chapters in this dissertation contribute methodologically to the extant literature by revising insights about measuring the inputs and outcomes of work. By observing differences between workers *within* the same occupation or the same level of education this dissertation unpacks the tasks that people do, the skills that people learn and the meaning they experience. This exercise of measuring labour inputs in a more disintegrated way does not solely contribute to an academic, methodological discourse: the conclusions of this dissertation can also guide policy in novel directions. In tackling the impacts of technological change on labour markets we need to understand the intricacies of how between-worker differences explain differences in labour market outcomes. The findings of this dissertation can inform the design of flexible, bottom-up and precise policy interventions that can help workers adapt within the existing structures on the labour market: by guiding workers in routine occupations towards performing nonroutine tasks, reshaping curricula in vocational education, or creating a work-environment with more possibilities to experience meaningfulness.

1.1 Setting the Scene

1.1.1 Technology, tasks and labour markets

This dissertation starts out from the literature on job polarisation. This literature focuses on the changing distribution of wages and employment on labour markets, and how that relates to technological progress and globalisation.¹ Job polarisation is

¹In this dissertation, I specifically focus on technology, and mostly those of the so-called 'third revolution': communication and information technologies – or computers (Brynjolfsson and McAfee, 2014). However, besides technology, globalisation has also greatly impacted labour markets over the past decades. Increased import competition following international trade (Autor, Dorn, and Hanson, 2013), and offshoring possibilities of jobs to other countries (Goos, Manning, and Salomons, 2014; Reijnders and de Vries, 2017; Terzidis and Ortega-Argilés, 2021) have reduced the demand for offshorable occupations

Figure 1.1: Polarisation of the Dutch labour market (1999-2012): Employment changes by occupational percentile

Source: Terzidis and Ortega-Argilés (2021). The skill percentile is defined by occupational mean wage in 1998.

the phenomenon where the number of jobs at the ‘tails’ of the income distribution are growing, whereas those in the middle are relatively declining. Compared to thirty years ago, relatively more people hold either high-skilled occupations – that require a college degree or higher – or service occupations – that involve assisting and caring for others. Manufacturing jobs and routine-intensive occupations are disappearing, either through offshoring or automation (Goos et al., 2014). This U-shaped pattern of employment growth has been documented in the US (Autor, Katz, and Kearney, 2006; Autor and Dorn, 2013) and many European countries (Goos, Manning, and Salomons, 2009; Goos et al., 2014). Relevant to this dissertation, it is documented specifically for the Netherlands as well, see Figure 1.1 (Ter Weel and Kok, 2013; Terzidis and Ortega-Argilés, 2021).

The phenomenon of job polarisation has been linked to the (non)routine task intensity of occupations (Autor and Dorn, 2013; Goos et al., 2014).² The notion behind this is that some technologies (such as computers) have a comparative advantage in

and jobs in import-competing industries.

²For clarification, tasks, jobs and occupations are three distinct elements. A task is a procedure that a worker executes. A job comprises the bundle of these tasks. An occupation is a statistical unit of measurement, where the average bundle of tasks performed by workers in that occupation is written in a dictionary of occupational titles. A ‘routine occupation’ is thus an occupation where the tasks description comprise relatively many routine procedures.

routine tasks, i.e. that they are relatively efficient in repetitive, rules-oriented and programmeable activities. When new technologies enter the market more tasks can be replaced, and this changes the allocation of tasks across workers and machines. This can be both in favour of employees – who can now use machines – but also against workers – when they see their tasks replaced. As ICT and other automation technologies became cheaper, they mainly took over many of the routine tasks. As a result, employment in these occupations declined. At the same time, humans have a comparative advantage in many other tasks, mainly those related to interpersonal contact, creativity or problem solving in unpredictable circumstances – also known as nonroutine tasks. Employment in nonroutine-intensive occupations grew, because these workers could be more productive by using technology. For instance, the accountant can let the computer do the bookkeeping, and focus on responding to clients’ more complex needs.

Most of the work that studies such task-based inequalities relies on occupational dictionaries, that describe the tasks and skill contents of each occupation.³ This type of data allows us to understand how the tasks people do at work relate to employment or wage growth in their occupations. However, relying on occupational task descriptions also requires the assumption that all workers in the same occupation perform the same tasks. Yet, there are a number of reasons why tasks may vary within occupations, especially in light of rapid technological change. First, and foremost, people have different abilities and preferences. We craft our own jobs (Wrzesniewski and Dutton, 2001), quite similar to students when collaborating in a group assignment. The one who is relatively skilled in Excel works on gathering the data, whereas the student with stronger presenting skills will likely present the findings. The task description (i.e. assignment) of these students would be the same, but the things they actually *do* are quite different.⁴ Second, workers can adapt to changing circumstances by switching to different tasks. For instance, when firms invest in technology, such as robots or computers, its workers start performing other tasks (Dauth, Findeisen, Suedekum, and Woessner, 2021). Spitz-Oener (2006) shows that occupations that are more exposed to computerisation show pronounced shifts away from routine tasks, and into more social and analytical tasks. However, as not every firm or industry adopts technologies at the same pace, the exposure to

³There are two main sources: O*NET and the DOT. O*NET (the Occupational Information Network) is an online database that lists the tasks, skills and broader characteristics such as work values and interests, for all occupations. The DOT (Dictionary of Occupational Titles) is its predecessor.

⁴This analogy also works for students working on an individual assignment: they focus on the parts of an assignment that they are comparatively good at, and their tasks depend on the tools they have at hand. For instance, some use $\LaTeX$ to write equations, whereas others write down their formulas by hand. The one task requires computer skills, whereas the other requires a steady hand.

technology may differ strongly across workers in different industries – which has consequences for the tasks they execute. Some office clerks have to work with paper-file cabinets, whereas others can make use of more advanced database systems.⁵ As such, their tasks and skill requirements are different, based on the relative adoption of technology in their environment.

This worker-level variation within occupations is important when assessing the risk of replacement of jobs by machines. One example where this became clear concerned the highly influential paper by Frey and Osborne (2017). By analysing both the tasks of occupations, as well as the scope for automation of each task, they concluded that 47% of jobs were susceptible to replacement by artificial intelligence. This number was picked up rapidly by various news outlets, sparking more fear of imminent mass unemployment due to technology – or a ‘robocalypse’ (Autor and Salomons, 2017). However, Arntz, Gregory, and Zierahn (2017) nuance the Frey and Osborne thesis: by studying how tasks vary within occupations, they show that the automation risk of (US) jobs decreases to 9%.⁶ This is still a significant number of jobs, but not as alarming. The crux of these differences lies in the underlying assumption of how people respond to automation technologies: the 47% is based on the idea that because a task can be automated, also will be automated – and that workers do not shift to new tasks. Arntz et al. (2017) show that it matters to measure tasks on an individual level: doing this, we can see that people that are more exposed to computerisation, actually perform different tasks. They adapt to their environment, which makes them less at risk of being replaced by technology.

Occupational data is useful for assessing the macro-level dynamics of task-based employment growth. However, to understand whether people performing certain tasks are doing better or worse in today’s labour market – i.e. to analyse which workers thrive – it is more informative to know whether a single worker performs more or less (non)routine tasks than their peers. In other words, one needs individual-level task data to understand the micro-level dynamics of routine-biased technological change, for instance how career growth and switching patterns relate to the routine-intensity of jobs. Furthermore, as tasks depend on the relative exposure to computerisation (Spitz-Oener, 2006), it is interesting to study differences in technology adoption for the understanding of how tasks shape inequalities.

⁵Even though there exist many technologies that can make work more efficient, there is always a lag in the adoption of it by *all* users. For instance, even though Microsoft stopped support for Windows 7 in January 2020 (after exiting mainstream support already in 2015), at that time still 25% of users had not yet upgraded to Windows 10. Source: <https://www.techspot.com/news/85480-windows-7-used-quarter-all-pcs.html>.

⁶By now, any citation of the Frey and Osborne (2017) paper is often accompanied by the Arntz et al. (2017) one.

These considerations formed the basis of my first project of the PhD: Chapter 2 of this dissertation. In order to study both within-occupation variation in tasks as well as between-industry variation in computer adoption, I use a large survey on Dutch workers. This survey, called NEA⁷, runs yearly since 2003, and asks respondents about their working conditions. A number of these questions were suitable for constructing a worker-level task index. This data is accessible via Netherlands Statistics: their system of microdata allows researchers to combine all sorts of data sets, through anonymous person-, firm- or school identifiers. As a result, I can not only study the results from the survey itself, but also link its respondents to their wages in the years following the survey, allowing a longitudinal study of the careers of Dutch workers, based on the tasks they execute at the time of the survey.

In Chapter 2, I find that routine task intensities can still explain pay differences between workers in this period, also when tasks are measured at the individual level. The Dutch labour market appears to be specifically characterised by positive returns to tasks that *complement* technology. Workers that perform relatively more abstract-oriented tasks – related to problem-solving, learning new things, and control over how tasks are executed – earn more, and this wage premium grows over the course of one’s career. This finding remains consistent over a number of specifications, regardless of the inclusion of, or interactions with, other variables. These findings fit into a simple theoretical framework in which nonroutine work complements ICT, and computers are routine-replacing.

Whereas nonroutine work is associated with higher wages across the entire labour market, the careers of routine-intensive workers are much more dependent on the environment in which they operate. Specifically, it matters whether the production process in an industry uses ICT intensively, and whether the industry is relatively dependent on routine or nonroutine tasks. When workers execute routine tasks in ICT-rich industries, this reduces their wages. This can be interpreted as a crowding out of routine work in industries where more tasks have been or can be replaced by computers. However, when a routine worker is surrounded by more routine workers, this positively affects wages: the return to routine tasks is positive in relatively routine-intensive industries and occupations. It is therefore not surprising that the results also show that when routine workers switch industries, they will switch to an industry that uses relatively more routine work than their initial industry.

One of the main contributions of this paper is the combination of individual-level task data from surveys and industry-level data on computer use. Though it is acknowledged that occupation-level data overlooks this within-occupation variation

⁷The *Nederlandse Enquête Arbeidsomstandigheden* – or Dutch working conditions survey.

(e.g. Autor and Handel, 2013), good quality individual-level data on tasks is not ubiquitously available, and knows many downsides of their own. For instance, surveys asking people about tasks are often cross-sectional, so one cannot track people and their tasks over time. Second, it is difficult to ascertain whether the survey questions provide good measures of the actual replaceability or complementarity of tasks. Third, many surveys have a limited number of observations per occupation, reducing the potential of picking up variation within narrow occupational classifications. And fourth, linking the task data to reliable information on earnings is not always possible, as most wage data from surveys originates from workers' own discretion in filling out the right numbers.

The data used in this chapter does not perfectly solve all these problems, but it does overcome a few. I use a large and nationally representative survey in the Netherlands to construct my task measures. The sample size (155,781 workers) allows for enough variation within occupations, and also between industries. As such, I am able to see how differences in tasks prices relate to the characteristics of the industry's production process: the average routine and nonroutine intensity, and the average use of computers. I link the survey data on tasks to register data on earnings and employment from Netherlands Statistics. Though the survey is cross-sectional (i.e. I can only observe the tasks and occupation of the worker at one point in time), the wage data are longitudinal: I can track the careers of workers over time. This data contains not only (tax-record) information on wages in the year of the survey, but also in the years after, plus information on the industry of employment, contracts, unemployment, working hours, and number of jobs. This additional information is one of the major advantages of the data used in this chapter.

Besides differences in pay, this chapter makes a case for studying the impact of technologies on labour markets through the channel of job quality. This is inspired by the work of Weil (2014). He argues that workplaces generate good and bad jobs, based on being on the payroll and receiving all the extra benefits of that, such as tenured contracts and lower risk of unemployment – besides higher pay. Good jobs are those in which people do not only have a sufficient hourly wage, but also sufficient working hours and job security. Bad jobs are those in which workers are in relatively more precarious working conditions, with all its consequences for their well-being (Cuyper and Witte, 2006). In this chapter, I show that this demarcation of good versus bad jobs can partly be explained by tasks. Routine work is associated with a higher probability of working in a non-tenured contract, being in an unemployment spell, having multiple jobs per year, working fewer hours and less often full-time, and lower job satisfaction in general.

1.1.2 Technology, skills and labour markets

Chapter 2 focuses on the tasks that people perform at their job, and how that relates to their careers. In Chapter 3, I move to the skills that people need to execute their tasks. The abilities we have are the main input that determine the success of the tasks we have to perform at work. The more complex a task is, the more complex the skills are that we need to execute this task. Figure 1.2 highlights this. First, observe the distributions of occupations along the horizontal (task) axis: workers in elementary level occupations have higher routine intensities and lower abstract intensities. Conversely, workers in more skilled occupations have higher levels of abstract tasks, and perform fewer routine tasks. Next, when we also observe the vertical (numeracy skill) axis, we see a monotonic relationship between cognitive skills and tasks: abstract tasks are associated with high numeracy skills, and vice versa for routine tasks.⁸

As the demand for tasks has changed over the past decades, the skill requirements have also changed. People either need different skills because they have to perform new tasks, or they need different skills to execute their ‘old’ tasks, but now using machines, such as computers. The question is: what skills do people need in today’s labour market and how do these change in light of technological progress? From Figure 1.2 we could make an inference: people with high cognitive skills work in abstract-intensive occupations, and those with lower numeracy scores are more often in relatively more routine-intensive occupations. Given that the demand for abstract tasks has increased since the onset of computers, a sensible deduction is that the demand for cognitive abilities has also increased over this period.

Indeed, when computers were first introduced in the workplace in the 80s and 90s, multiple authors documented a sharp rise in the wages earned by university graduates (Katz and Murphy, 1992; Goldin and Katz, 2010). This contributed to a theory of *skill*-biased technological change: technology works well with high-skilled workers, which has drastically increased the average wage of those with a college degree. Especially computer-intensive industries now employ relatively more college-educated workers than they did historically (Autor, Levy, and Murnane, 2003). The theory of *task*-biased technological change nuanced this view by highlighting that service-oriented jobs in the lower skilled part of the labour market have also seen

⁸The data used for Figure 1.2 is from PIAAC, the Programme of the International Assessment of Adult Competencies. It tests adults on their skills in numeracy, literacy and problem solving in an IT environment. Each respondent also filled out an extensive survey of background questions, including on their occupation and the tasks they performed. For my master thesis, I used this data set to construct worker-level measures of tasks and analysed the returns to tasks in an international setting. The data also lend itself well for figures in an introduction.

Figure 1.2: Correlation of numeracy skills and task intensities on the occupation level

Note: Author's calculations using PIAAC 2012 and 2017 data. The abstract and routine task intensity are constructed using a principal component analysis on survey questions related to tasks, equal to the ones used by Rica, Gortazar, and Lewandowski (2020). These indices are standardised (with mean zero, and a standard deviation of one) across the population. Both the task indices and the numeracy skills are averaged over 2 digit ISCO occupations. Black line is fitted regression.

growth because of their nonroutine intensity (Autor and Dorn, 2013). Nevertheless, even though nonroutine low-skilled occupations grew, many low and middle-skilled jobs have disappeared, because these are the types of jobs that use routine tasks intensively (Autor and Dorn, 2013). In other words, technology is biased against low- and middle-skilled workers, conditionally on that these workers execute routine-intensive tasks.

Despite all this, in the paper *Why are there still so many jobs?* David Autor (2015) predicts that “a significant stratum of middle-skill jobs combining specific vocational skills with foundational middle-skills levels of literacy, numeracy, adaptability, problem solving, and common sense will persist in coming decades” (p.27). The solution to making sure that workers in the future will have the right skills will likely lie in schooling, and making sure graduates have the right skills upon entering the labour market. This makes the education system both the solution, as well as the bottleneck. Autor specifically mentions this as the main catch of his prediction:

The ability of the education and job training system (both public and private) to produce the kind of workers who will thrive in these middle-

skill jobs of the future can be called into question. In this and other ways, the issue is not that middle-class workers are doomed by automation and technology, but instead that human capital investment must be at the heart of any long-term strategy for producing skill that are complemented by rather than substituted for by technological change. (Autor, 2015, p.27)

Even though the decline of middle-skilled jobs is widely documented, and the importance of education for the middle-skilled is acknowledged, the majority of the literature on skills is focused on college-educated workers.⁹ Unfortunately, this means we know relatively little about which skills should be part of the middle education system in order to educate thriving workers – while still be most at risk for automation. This preoccupation is not only seen in academia: in politics, the middle-skilled education sector also receives far less attention than higher education (Schakel and Van der Pas, 2021). Sparked by the above-mentioned quote, I wanted to dive deeper into the position of the middle-educated worker.¹⁰

Two questions arose. First: how are the middle-educated currently doing on average? And second, who is thriving and who is not? To provide some data for answering the first question, Figure 1.3a and 1.3b sketch the picture. In the top panel, we see that the wage distributions of *both* middle and lower educated workers are skewed to the left towards the lower wage percentiles, whereas the wage distribution of higher educated workers is strongly centered around the top earning jobs. In other words, it seems that higher educated workers have access to high-paying jobs, whereas middle-educated workers are relatively more often competing with lower-educated workers for the jobs in the lower ends of the wage distribution.¹¹ However, Figure 1.3b suggests that these differences cannot be completely explained by the level of cognitive skills: higher educated workers do not score disproportionately better in numeracy, literacy and problem solving tests. Moreover, we see that there is substantial variety in wages between middle-educated workers: some workers are doing just as good as higher educated workers, whereas others are in the same wage groups as lower educated workers (notably, this last group is larger). This variation

⁹Novel work by Acton (2021) studies the effect of local labour market recessions on the student's choice of entering different fields in community colleges in the US, the equivalent of the Dutch mbo. She also highlights this gap in the literature on student choice: most of it concerns 4-year college degrees.

¹⁰When I mention higher education, I refer to the International Standard Classification of Education (ISCED) levels 5 and 6 (Dutch hbo and university). Middle education refers to ISCED 3 and 4 (Dutch mbo level 2 to 4 or a havo/vwo degree). Low education is ISCED 1 and 2 (Dutch mbo level 1, vmbo degree or no education).

¹¹This finding also remained significant in a series of estimations that include a full set of demographic controls: the return to higher education is significantly higher than that to middle education, and, importantly, the difference between high and middle education is significantly *larger* than the difference between middle and lower educated workers.

Figure 1.3: Wages and skills of low, middle and higher-educated workers**(a) Wage distribution** **(b) Cognitive skills distribution**

Note: Author's calculations using PIAAC 2012 and 2017 data.

can be explained by differences in the general abilities or the field of education: students majoring in STEM-degrees – related to Science, Technology, Engineering and Math – tend to have higher wages than those with service-oriented degrees (Altonji, Blom, and Meghir, 2012).

The second question then becomes: which middle-educated workers are doing well in terms of wages, and which are not? There is evidence that middle-skilled workers switching out of routine occupations, and into more cognitive occupations are doing relatively well. However, those that stayed in routine occupations, for instance because they did not have the right skills to move to other types of jobs, have slower wage growth (Cortes, 2015). The literature also shows that mostly the *young* middle-educated workers are doing poorly. Recent research by Ter Weel, Zwetsloot, and Bisschop (2021) describes how technological progress has worsened the labour

market outcomes of Dutch mbo graduates. Specifically, those students educated in degrees preparing for relatively routine-intensive occupations had significantly lower employment probabilities. Similarly, Reinhold and Thomsen (2017) document the ‘declining fortunes’ of young workers in Germany, and particularly of the middle- and low-educated workers. This may have to do with technology adoption (Dauth et al., 2021): when labour-replacing industrial robots are adopted in local labour markets, incumbent workers are assigned new tasks. However, young workers that were just entering the labour market faced lower labour demand. As a result, many of the middle-educated workers adapted by obtaining college degrees or were forced to switch to other jobs – which then faced higher competition.¹²

A few questions kept puzzling me during the first years of my PhD: which of the middle-educated students are currently doing better, and does that depend on what they learned in their classrooms? And what can be changed *within* the Dutch mbo to prepare these students better, rather than only facilitating the road to higher education? At this point, it might be worth mentioning that simultaneously with the PhD program, I was also co-founder of a start-up that focused on increasing student engagement in the Dutch vocational education system.¹³ Initially, these were two completely separate worlds: I focused on the study of routine tasks in the PhD for four days a week, and one day a week I spent my time in varying mbo schools.

While working in the vocational education system, I learned about the existence of ‘kwalificatiedossiers’: the qualification files that describe the skills students should have learned upon graduating. I realised that this could be an interesting source of information, for a number of reasons. First, these files were all set up in a similar way. All degrees described the skill requirements of their graduates in identical ways, facilitating data collection. Second, the files are constructed on a national level, implying that a carpenter in Groningen will learn the same skills as a carpenter in Rotterdam. And third, the Netherlands has a strong accountability system: the Dutch Inspectorate of Education regularly checks whether the examination complies with the qualification files. In other words: the structure of the files makes data harmonisation possible, the set-up of the system warrants the exogenous determination of the skills that are ‘supplied’ to students, and the files are a trustworthy source of information on what is actually going on in schools.

¹²The long term result might be positive, as students making the choice to acquire more education could potentially be better prepared for the future labour market. However, not every student will be able to do so.

¹³This was a spin-off project of the *Nationale DenkTank* (National Thinktank), in which I participated in 2016. See www.nationale-denktank.nl/jaarlijkse-denktank/denktank-2016 for more information. The theme of 2016 centered around preparing the Dutch vocational education system for the future (labour market).

From these qualification files, I was able to extract novel skill measures for all the graduates of the Dutch education system. Of course, I could have expected that this task was more cumbersome and time consuming than anticipated.¹⁴ With this new data set of degree-level skills for all middle-educated graduates in the Netherlands, I can go beyond the major and field of education. This kind of data helps in our understanding of which types of skills students learn – and how that predicts wages in the first years after graduation.¹⁵

Chapter 3 is based on this data, and it presents three main findings. First, graduating from relatively social-skill intensive degrees is negatively associated with wages in the first years after graduation while technical skills are associated with positive returns. Both relations persist until at least 10 years after graduation. Second, demand for technical, social and basic skills differs strongly across fields, levels and tracks of education. For instance, students that graduate in the health-related field of education have higher returns to technical skills, as compared to STEM and economics graduates. Third, wage returns to skill are conditional on the sector of employment: social skills are more strongly negatively associated with wages in the high skill service sector than in the low skill service sector. Interpreting these results, we might argue that degrees focusing relatively more on social skills have lower wage returns. Importantly, this does not necessarily imply that the actual demand for social skills is lower for middle-educated students – but it might be more worthwhile to *teach* other skills in school. This could have potential implications for the construction of curricula.

1.1.3 Technology, meaning and labour markets

Though analysed in-depth in psychological, philosophical, and organisational research, meaningful work is not a prominent theme in the literature on technological change and labour markets – or modern economics in general. However, the study of work as a source of meaning could bring a valuable and important contribution to our field. When both the contents and the skill requirements of our jobs are subject to drastic changes, how can the meaning we derive from our work go unaffected?

In economic models, supplying labour is perceived a disutility: an unpleasant activity that must be endured as a means to earn an income and finance consump-

¹⁴I used natural language processing programs called *Frog* do to this (Van den Bosch, Busser, Canisius, and Daelemans, 2007). I owe eternal thanks to our colleagues at the CIT service desk for their help in getting *Frog* leaping.

¹⁵I realise that this focus on wages is very narrow and overlooks other reasons why a skill should or should not be taught. Therefore, in the chapter I aimed not to make normative statements on the in- or exclusion of skills in curricula.

tion.¹⁶ However, self-determination theory from the psychological literature focuses on how intrinsic motivation makes people willing to work *irrespective of pay* – and can even derive utility from working (Ryan and Deci, 2000a). It states that once people have satisfied three basic psychological needs, they will feel motivated and exert effort. The three needs are *competence* (the experience of mastery), *autonomy* (the feeling of choice and control), and *relatedness* (the feeling of being connected with and belonging to other people). Variation in the satisfaction of these needs explain why some people exert more effort than others. This type of research calls for organisational behaviour that facilitates the satisfaction of basic psychological needs in order to ensure motivated workers, and thus, in more economic terms, productive workers (Ryan and Deci, 2000a)

In order to experience meaning from work, the three psychological needs have to be satisfied. Through autonomy, people feel that they have made their own decisions in doing their work, rather than e.g. following a machine or strict orders. As such, autonomy can make work more fulfilling. When people experience that they are skilled at what they do, they feel that their own personal actions contribute to the final output of a task: their competence matters for the success of a task. And lastly, the feeling of being appreciated and supported by co-workers or supervisors i.e. relatedness, also increases the meaning we derive from our tasks. People think of their tasks as more fulfilling when others do too. Motivation and the experience of meaning are strongly interrelated, and this matters when the skills and tasks of people change.

To understand the channels through which automation can affect meaning and motivation, one can refer to case studies on technology adoption in the workplace. One interesting case is provided by Barrett, Oborn, Orlikowski, and Yates (2011), who study the impact of a drug-dispensing robot in a pharmacy. They show how the robot affects the meaningfulness of three different occupational groups: pharmacists, assistants and technicians. First, pharmacists indicated that the increased delivery speed of medication had improved their job, by creating more time for in-depth counseling to patients. This made their work more interesting, and appealed stronger to their sense of competence. The assistants to the pharmacist had opposite experiences. Originally working independently, they now had to follow the working pace and tasks of the robot, over which they had little control. Their sense of competence also decreased, as the original expertise of knowing where to put which medicine

¹⁶Labour supply models are also founded on the assumption that people will maximise the monetary returns to labour. However, people do not solely choose to work in order to maximise wages. This has resulted in models of compensating differentials, in which workers might ask a premium to perform undesirable tasks, or accept a lower wage for desirable ones.

was no longer necessary. The third group, the technicians, had yet again different experiences. Before the robot, they had similar responsibilities as the assistants. The introduction of the robot changed this. As the robot often stagnated and the technicians were the only workers authorised to fix the problems (even if the assistants knew how to), their sense of competence and status within the organisation increased (Barrett et al., 2011). Thus, as work changes, the experience of meaningfulness also changes. However, this relation is not straightforward: each occupational group may be affected differently.

Beyond a fear that jobs disappear and wage inequalities aggravate when technology progresses, there is an equally understandable fear that the ‘remaining’ jobs will be demeaning (Shiller, 2019). As presented in the example above, this fear may partly ring truth, but might as well go in the opposite direction. On the one hand, the hopeful thing about technology is that it mostly substitutes the relatively mundane tasks, that give little rise to either meaning or self-determination. If the task substitution is accompanied by new tasks, this may result in a higher experienced level of meaning. However, the fear of demeaning work might be legitimate if only very specific tasks are left for human labour, that create little to no variation, challenge or interesting work (Smids, Nyholm, and Berkers, 2020).

Chapter 4 aims to stimulate economists to take meaningful work into account when studying labour markets. I do not think the intrinsic value of meaningfulness is disputed: work is an important source of meaning that strongly relates to the identity of people and their sense of fulfillment in life. However, its economic relevance is not always clear. The chapter shows that meaningful work is in fact related to worker behaviour, such as sick leave and participation in skills training, which should be of interest to economists. Furthermore, the chapter also makes a theoretical contribution, by combining motivation theory from the psychology literature to an economic labour supply model. In contrast to most labour supply models, I propose a framework in which motivation is a precondition for supplying labour, and where people can obtain utility from doing meaningful work.

1.2 The future of (my) work

Finally, I would like to take a look forward: what I see happening concerning the future of work, the future study of work, and my own work.

To begin, the future of work will center around two main questions: to what extent will we adopt technology, and how do we adapt? The first one will be a result of the (political) choices we make today. How technologically advanced do we want

our labour markets to be? This will depend on actors in the public and private sector to shape a future where we can reap the profits of technological progress. Ideally, these choices would make sure that profits are distributed in a fair way. Too strong inequalities may possibly undermine democratic institutions, and should therefore be of interest to not only researchers and politicians, but also to the firms that adopt technologies.¹⁷

The second question is focused on how we live with technological change, once it is there. Whether we are able to adapt to changing circumstances will depend on our own qualities and possibilities. The human nature is adaptive, and those who will fit best to the new technological world will most likely do better. However, there is fear that ‘this time it’s different’, as self-learning machines are capable of more complex things than we can possibly imagine. Though it is true that artificial intelligence will go beyond many technologies we have seen before, that does not mean we should sit idle and wait for technology to replace our tasks. Rather, it calls for a proactive attitude, not only in ourselves, but also in our institutions. It begs for the question how we can prepare students for a labour market that continuously changes, and how we can ensure that flexibility and resilience become core characteristics of the new cohorts of graduates.

In this context, research plays an important and informative, role, in which I hope to make my own modest contribution. I envision (my) future research in this topic to be mainly centered around i) the role of tasks and skills in labour market inequalities and ii) conceptual and empirical work on a broader economic perspective on human thriving.

One of the main contributions of the first two chapters in this dissertation is through the channel of data. In Chapter 2, I generated a new data set for the study of individual-level tasks in a longitudinal wage setting, by making optimal use of the many possibilities that register data from national statistical agencies provide. The access to such data has drastically improved over the past years, and the ease with which high quality data can now be combined and analysed is unprecedented. This has made way for the study of new questions, that could not have been studied before. For Chapter 3, I generated a new data set on skills using recent innovations in Dutch language processing. I deem especially these innovations in analyzing text data to be a promising future avenue of research. There is an abundance of text to analyze, which allows us to understand economic mechanisms from yet untapped

¹⁷The consequences of technological change and job polarisation are not economic in nature alone, they also relate to e.g. electoral results (Kurer and Palier, 2019) and marriage and divorce rates (Greenwood, Guner, Kocharkov, and Santos, 2016).

sources. Besides this chapter, similar works have been published (e.g. Frey and Osborne, 2017; Webb, 2019), and I hope many will follow in taking advantage of the complementarity between artificial intelligence and (economic) research tasks.

Even though the literature on job polarisation is well-established, its conclusions are not unchallenged. One of the main concerns is the literature's strong focus on occupational employment trends, which are used as indicators of changes in job skill requirements. However, as many trends occur *within* occupations, the complete picture might be different. For instance, focusing on wage bins, rather than average wages and employment within occupations, Hunt and Nunn (2019) show that employment growth was mostly concentrated in high-wage jobs versus, essentially, all other jobs: this way of measuring provides no evidence for increases in employment for low-wage jobs. They argue that the U-shaped pattern, and specifically the growth in the lower end of the occupational wage distribution, might be an artifact of occupational coding. Likewise, the Frey and Osborne (2017) versus Arntz et al. (2017) example highlights that occupational trends may not be informative of actual changes, as it overlooks this within-occupation variation. Studying such micro-level dynamics, and especially those within occupations, might be relevant for future research aiming to understand the intricate processes underlying job polarisation.

These concerns need not do any injustice to the interpretation that some skills work better with technologies, and that inequalities arise because of variations in skill-technology complementarity. However, it might also be worth considering that the distinction between routine and nonroutine tasks might proxy for other micro-level dynamics than merely abilities to work with or against machines. For instance, Weil (2014) documents that firms increasingly center their organisations around core tasks, where they outsource peripheral tasks to lower-tier firms (both domestic and foreign). This process, described as 'workplace fissuring', creates task-based inequalities in both job security and wages. Arguably, routine and nonroutine tasks could play on this similar spectrum of core versus peripheral tasks, or offshorable versus location-constrained tasks. On a fundamental level, the distinction between routine and nonroutine tasks might go deeper than what workers *do*, but centers around concepts of responsibility, autonomy and bargaining positions of workers. Such a discussion might clear the way for task-based explanations for growing inequalities, that are not founded in a technology-centered story. The task approach can then also be used in existing theories that aim to explain growing wage inequalities, such as organisational change (Caroli and Van Reenen, 2001), offshoring (Grossman and Rossi-Hansberg, 2008; Baumgarten, Geishecker, and Görg, 2013), structural change (Bárány and Siegel,

2018), or growing female labour market participation (Wright and Dwyer, 2003). Through linking tasks to worker power and autonomy, inequalities may even relate to labour market institutions, such as de-unionisation (Card, Lemieux, and Riddell, 2018) or the minimum wage (Salverda and Checchi, 2015).

Lastly, I see a way forward in labour economics that focuses on human thriving in the workplace in a broader sense than wages and employment alone. Following this work, Chapter 2 makes a case for using data on other job indicators besides wages. This allows us to paint a more complete picture of career inequalities between workers, based on the tasks they execute. Besides job quality, I believe the constructs of meaningful work and motivation may become more relevant in the coming period, especially in light of technological progress. It is likely that technologies generate inequalities in the possibility of jobs to foster human flourishing. In economics, this is a relatively under-researched topic. In part, this may be explained by a difficulty to conceptualise the channel through which inequality in meaning will materialize. As reflected by the pharmacist-example presented above, it seems that the effect of technology on meaning will depend on the tasks workers currently execute, and the scope for task substitution, creation and augmentation. Also, it will depend on our individual abilities to craft jobs in response to new technologies, and the willingness of employers to think along the well-being of their employees. Conceptualizing this into a framework is not straightforward. Moreover, a comprehensive empirical analysis will require more data that combines worker-level tasks, technology adoption and workplace characteristics. This makes it a challenging new avenue of research. Yet, given the importance of work for human well-being, it is too important to be left to case studies alone. It deserves a wider audience – and one that includes economists.

Table 1.1: Thesis structure

	Chapter 2		Chapter 3	Chapter 4
Theme	Tasks		Skills	Meaning
Article status	Working paper		GLO Discussion paper, Revise & Resubmit at an international journal	Published in Labour Economics
Main topic	Wage inequalities based on worker-level tasks and industry-level computer adoption		Early career wage inequalities based on skills taught in vocational curricula	What makes work meaningful and how it relates to labour market behaviour
Main independent variable(s)	Routine and nonroutine abstract tasks		Skill frequencies (social, technical and basic cognitive)	i) Psychological needs (autonomy, competence and relatedness), ii) meaningfulness
Main dependent variable(s)	Hourly wage in year of survey, wage growth in years after survey, industry-switching patterns, job quality indicators		Hourly wage in first year(s) after graduation	i) Experiences of meaningfulness, ii) sick leave, retirement decision, training participation
Data sources	Dutch Working Conditions Survey (NEA), CBS Microdata, EWCS		Text data extracted from S-BB Qualification files, CBS Microdata	European Working Conditions Survey (EWCS)
Time period	2008-2019		2010-2019	2005, 2010, 2015
Sample size	155,781 Dutch workers + 53,186 European workers		322,205 Dutch graduates	48,420 European workers

2

Technology, Tasks and Careers

2.1 Introduction

The technological advances of the past decades, and specifically computers, gave rise to efficiency increases in the execution of programmable, rules-oriented activities: routine tasks. This development has benefited those executing abstract, problem-solving and creative tasks, as their productivity increased when working with computers. It is, however, to the detriment of those in routine-intensive occupations, who have seen their tasks replaced. The routine-replacing nature of computers has had serious consequences for the demand for routine labour. First, employment in routine-intensive occupations has declined, accompanied by a relative decrease in wages compared to nonroutine workers (Autor et al., 2003; Autor and Dorn, 2013; Goos et al., 2014). Second, workers within all occupations are gradually moving towards the execution of nonroutine, abstract tasks, and especially when exposed to computerisation (Spitz-Oener, 2006; Arntz et al., 2017). This has thus reduced the average routine-intensity of jobs on western labour markets.

Nevertheless, technology has not fully eradicated routine work. Whether all routine tasks in a production process will be replaced depends on i) the scope for technology adoption: the suitability of the tasks for replacement and the willingness of firms to invest, and ii) the scope for worker adaptation: the ability of workers to change their tasks following organisational change (Walk and Handy, 2018; Tims and Bakker, 2010). This heterogeneity in adoption and adaptation can lead to variation in tasks between workers in the same occupation, but employed in different sectors (Bárány and Siegel, 2020). Tasks might or might not be replaced, and workers might or might not adapt their tasks when exposed to new technology. Given this, occupation-level task data might not provide a ‘true’ image of the set of replaceable tasks executed by workers, as it overlooks the changing distribution of tasks across industries and skill-levels. Measurement error in tasks may obfuscate the direct relationship between tasks and wages, and as such hinders the study of individual-level dynamics of routine-based technological change (RBTC).

In this paper, I analyse the extent to which RBTC still describes patterns of wage inequality between workers after 2000, given that many routine-replacing technologies have been adopted over the past decades where possible and workers have been

This chapter has benefited from the help and support from Robert Inklaar, Milena Nikolova and Steven Brakman. I would also like to thank Ronald Bachmann, Gaaitzen de Vries, Marcel Timmer, Maria Balgova, Terry Gregory, Joana Duran Franch, Melline Somers, Pascual Restrepo, Anna Salomons, Alexandra Spitz-Oener, two anonymous referees and participants at EALE (2020), EEA (2020), LISER-IAB (2020), ESPE (2019), PEGDECH seminar (2019), LEER (2019), SOM PhD Conference, ROA workshop (2018) for helpful comments and advice. Part of this paper was revised during a research visit at IZA in 2020. Any remaining errors are my own. Personal contribution to this chapter: single-authored.

able to adapt accordingly. More precisely, I study the earnings and employment dynamics of workers conditional on their tasks and the computer adoption in their industry. The novel aspect of this paper is the use of individual-level task data from a nationally-representative working conditions survey, combined with non-public microdata on labour market outcomes from Statistics Netherlands. The data allows me to move beyond occupation-level descriptions of tasks, and observe within-occupation and between-industry variation in tasks. These differences between workers in the same profession, either through different exposure to technology in their industry or variations in the scope for adaptation, may be informative of explaining wage inequalities on the current labour market.

Figure 2.1 paints a picture of between industry variation in tasks on the Dutch labour market. I have plotted the ICT-intensity of 2 digit industries (measured as average computer use in an industry) against both the average routine (Figure 1a) and abstract (Figure 1b) intensity of workers in that industry. These means are constructed using the individual-level task data. There is a clear pattern of higher abstract-intensity of workers in ICT-intensive industries, versus a monotonic negative relationship between routine tasks and increasing computer use in industries. In terms of an industry-level production function, this reflects the substitutability of routine work and computer capital. The pattern is thus also potentially indicative of RBTC. Routine work has been replaced by computers, and the result is relatively more abstract work in these industries than in those with lower ICT.

This paper uses these variations in tasks and ICT to explain wage differences and employment dynamics between workers. The paper makes three main contributions. First, I outline a simple theoretical framework that allows for within-occupation variation in tasks and between-industry variation in ICT adoption, in order to derive micro-level conclusions on RBTC-induced wage and employment inequalities. Second, I create a routine and nonroutine abstract index using individual-level task data, based on a Dutch working conditions survey (WCS). The underlying elements in these indices also appear in similar fashion in the American and European WCS and can thus be used to study similar dynamics in other countries.¹ Third, I test the conclusions from the framework using a linked employer-employee data set, in which the survey data are merged to rich administrative data on earnings and employment from Statistics Netherlands. In the data, I observe workers not only in the year of the survey, but also in a period of up to 8 years thereafter. I can therefore show how current tasks influence current wages, but also the later growth paths and

¹I explore this in the appendix of this chapter, where I rerun some of the analyses using the European working conditions survey (EWCS).

Figure 2.1: Smoothed means of routine and abstract-intensity by standardised ICT intensities on the industry level.

A. Smoothed mean of routine intensity over industry ICT intensity

B. Smoothed mean of nonroutine (abstract) intensity over industry ICT intensity

Note: Dots represent employment-weighted industry means in routine and abstract-intensity. Task intensities are sampling-weighted standardised values across entire population by year of survey. ICT-intensity is calculated as the industry mean of workers' daily computer use in hours, standardised for the entire sample. Source: Author's calculations using non-public microdata from Statistics Netherlands.

switching patterns of workers.

My findings show that the Dutch labour market can still be characterised by routine-biased technical change in the period between 2008 and 2019. I present the following findings that support this conclusion. First, I show that a one standard deviation increase in performing nonroutine abstract tasks is associated with higher hourly wages of 4.1%, even when comparing workers within the same 4 digit occupation. Second, a longitudinal analysis shows that these inequalities aggravate over time. After 5 years, the growth rate of nonroutine workers is .47 percentage points higher. Furthermore, abstract and routine tasks are also significantly related to various elements of job quality, such as contract tenure, job satisfaction and contracted hours. This highlights that tasks are not only shaping inequalities in wages, but also in broader job quality dimensions.

To provide evidence for the claim that these patterns can be explained by RBTC mechanisms, I account for the possibility that a higher return on abstract-intensity is capturing general ability, i.e. abstract-intensive tasks are simply more complex. In that case, we would see that nonroutine and routine intensities should have the same return across industries, regardless of the average routine or abstract-intensity of the production process in that industry. My results show that routine work has higher returns in industries that use routine tasks intensively. This finding shows potential selection on comparative advantages rather than absolute advantages in tasks, which would be the case where (non)routine workers earn (more) less in all industries because of (higher) lower abilities. As a second test for RBTC mechanisms, I interact worker-level tasks with industry-level measures of ICT use. This is a direct test of complementarity or substitution between tasks and technology. Consistent with an RBTC framework, routine tasks are associated with even lower wages if workers execute them in ICT-intensive industries.

Lastly, I show that the (non)routine-intensity of tasks is associated with different switching patterns between industries. If workers are relatively more routine-intensive than their peers in the same industry, they have a higher likelihood of switching to a more routine-intensive industry. On the other hand, relatively abstract-intensive workers are less likely to switch to a more routine-intensive industry. They also do not tend to move to industries with either lower abstract-intensity or lower ICT-intensity as compared to the industry where they initially were employed. These switching patterns show that workers sort according to comparative advantages. Workers with higher efficiency in a certain task move to industries that use that task intensively. Following the routine bias in technology, I show that industries that use relatively more ICT than the 'home' industry of the workers attract abstract-intensive

workers, whereas routine-intensive workers are more likely to move to industries that use less ICT than their current industry.

2.2 Related literature

2.2.1 Routine biased technological change

By now, it is widely accepted that technical change has been more favorable to nonroutine workers, and detrimental for routine workers. Most of this research is based on the use of occupation-level task data. Using such data, research has linked the execution of routine tasks to job and wage polarisation, as most jobs in the middle of the wage distribution tend to contain routine tasks. This resulted in a U-shaped pattern of growing employment and earnings in the tails of the income distributions in the US (Autor and Dorn, 2013) and many European countries (Goos et al., 2014), among which the Netherlands (Van den Berge and Ter Weel, 2015). Routine-biased technological change creates a dichotomy in the labour market between high-skilled, high-wage work ('lovely jobs') and low-skill, low-wage work ('lousy jobs'), with little in between (Goos and Manning, 2007). And though employment for low-skilled work has increased, this is not accompanied by the same wage growth that has become apparent for high-skilled workers. The creation of jobs at the bottom does therefore not outweigh the disappearing of jobs in the middle, especially in terms of job quality.

Employment polarisation is hindering the possibility of low and middle-skilled workers to access qualitatively good careers when middle-skilled jobs as stepping stones are disappearing. At the individual level, there is evidence that workers in routine occupations have relatively lower wages compared to workers in more abstract-oriented occupations (Cortes, 2015; Cavaglia and Etheridge, 2017; Böhm, 2020). Furthermore, workers in routine or otherwise declining occupations also face worse career trajectories in the medium to long run (Edin, Evans, Graetz, Hernnäs, and Michaels, 2019). Unless workers are able to make a switch to nonroutine occupations, which mostly only the relatively skilled-workers in routine occupations actually can, routine workers have less favorable careers (Cortes, 2015).

Job polarisation can also be explored using other indicators of job quality beyond wages, rather than wages. Weil (2014) demonstrates that as firms are forced to specialize, they focus more on their core competencies. As a result, they tend to shed work that does not directly create value. Though he does not immediately relate this to routinisation, the obvious consequence is that abstract workers are more likely to

be a part of the core business than routine workers, being that they are able to create value for the company. Displaced workers end up in subcontracted or outsourced contracts, where they have both lower wages, less extra benefits and less job security. This finding will motivate an analysis of job quality, where I show that abstract work is related to more stable employment in terms of tenure, full-time work, number of jobs per year and unemployment, but also more job satisfaction. The converse is true for routine work. The findings in this paper thus highlight that job stability and quality might be an important, yet under-researched element of technological progress. They also naturally follow from the model as job quality indicators can be seen as an element of the wage bill: providing tenure is costly to the firm, even though it is a non-pecuniary outcome for the worker.

The above-mentioned studies on micro-level implications of RBTC use occupation-level task data. The present paper fits in the literature that exploits within-occupation variation in tasks. For example, Marcolin, Miroudot, and Squicciarini (2016) create a measure of tasks on the individual level. Besides showing that within-occupation variation is substantial, they also show that routine-intensity varies substantially across countries and tends to be (weakly) associated with lower education as well as lower skills in literacy, numeracy and problem solving. Other research using individual level tasks also highlight the existence of within-occupation variation. Autor and Handel (2013), Arntz et al. (2017) and Stinebrickner, Stinebrickner, and Sullivan (2018) show that tasks substantially vary within occupations, and Arntz et al. (2017) specifically mention the increased variation in occupations that are more exposed to task replacement. Furthermore, using several waves of a German Qualification and Career Survey (GQCS), Spitz-Oener (2006) shows that occupations are also changing over time, and increase in skill requirements. This is especially relevant for rapidly computerising occupations. Moreover, Cassidy (2017) and Akçomak, Kok, and Rojas-Romagosa (2016) demonstrate that these changes over time occur both within and between occupations. For the period after 2000, Lewandowski, Park, Hardy, and Du (2019) use survey data from across the world to show cross-country differences in tasks, and they provide evidence for a shift away from routine to nonroutine abstract tasks, which is strongly related to a country's economic and technological development.

A sub-branch of this literature aims to estimate the returns to tasks defined on the worker-level. For instance, Autor and Handel (2013) and Cassidy (2017) show that tasks explain a significant portion of variation in wages in the US and Germany respectively, even when accounting for occupations up to a 3-digit level and demographic characteristics. Cassidy (2017) uses two waves (1986 and 1992) of the

GQCS, whereas Autor and Handel (2013) rely on the Princeton Data Improvement Initiative (PDII) survey. Most recently, Rica et al. (2020) use PIAAC data to show that nonroutine task premia are higher in countries with lower ICT levels and more labour market protection. The finding on the complementarity between ICT and nonroutine tasks is also tested in this paper, though I find less evidence for the complementarity between nonroutine work and ICT-intensity, and more so for the substitution of routine tasks in the presence of ICT. Ross (2017) uses a synthetic panel of tasks, by linking individual data to several waves of O*NET data, thereby circumventing the need for survey data. He finds that the effect of routine-intensity on individual wages in 2013 is lower than in 2004, and concludes that it decreases over time.

However, the cross-sectional nature of most survey data does not allow for longitudinal wage analysis. Wage differences emerging due to sorting into different tasks at an earlier stage can thus not be explained. This paper aims to fill this gap, by adding longitudinal register data on earnings and employment to cross-sectional survey data. To my knowledge, the only paper that uses survey data to relate individual tasks to future wages is that of Stinebrickner et al. (2018). Although they are mainly concerned with human capital accumulation through the execution of complex tasks in the years after college, they describe the role of complexity tasks in the determination of both current and future wages. As is to be expected, workers performing more complex task have higher wages and wage growth.

2.2.2 Industry-level technology adoption and worker-level adaptation

A micro-level literature explores the the consequences of technology adoption and worker productivity (Entorf, Gollac, and Kramarz, 1999; Borghans and ter Weel, 2007; Bresnahan, Brynjolfsson, and Hitt, 2002, e.g.). This scholarship finds that workers' productivity change when they are exposed to certain technologies at the workplace. This allows for the direct estimation of the effect of, for instance, computer use on productivity. It is clear that there are productivity gains, especially for highly-skilled workers. As a result, there is also self-selection based on skills in terms of who make a decision of switching to with computers and who do not, such that the actual productivity increase is smaller than the difference between computer users and non-users (Entorf et al., 1999; Borghans and ter Weel, 2007). On the firm-level, technology adoption is also associated with productivity increases. Clear complementarities emerge between IT adoption, workplace reorganisation and hiring of more labour that can use technologies, and the creation of new products and services (Bresnahan et al., 2002).

When studying the relation between (routine) tasks, technical change and micro-

level outcomes, one should explicitly conceptualise how the technology of interest operates in the workplace. Each type of technology changes the organisation of work in its own distinctive way. For instance, Webb (2019) shows how software, robots and artificial intelligence each replace different tasks by analysing the overlap between patent texts and occupational descriptions. The tasks that can be executed using software correlate most strongly with routine cognitive, routine manual and nonroutine manual and least with nonroutine cognitive tasks. Because I specifically focus on computer technologies and thus software, I use this finding as empirical basis for the creation of two task-based measures. My measure of nonroutine task-intensity aims to capture nonroutine cognitive tasks elements, and focuses on the importance of learning, problem-solving and task discretion.² The measure of routine-intensity captures routine cognitive and routine manual elements. As inputs, I use data on repetitive movements, and autonomy in the speed and sequence of tasks. These questions reflect the rules-oriented, procedural element of routine tasks, as well as the repetitiveness in motion accompanying routine work. Aggregating my task-based measures to occupations, they highly overlap with occupation scores for routine and abstract-intensity from Autor and Dorn (2013), implying they capture the same underlying description of tasks.

2.3 Conceptual framework

The aim of this paper is to understand how technology affects different groups of workers on the labour market. Following the literature on routine-biased technical change, a large share of the variation in wages and career development should be explained by the tasks that workers execute. In this section, I develop a simple model that explains how routine and nonroutine tasks shape wage inequalities between workers, where I specifically allow for within-occupation variation in tasks. The novel element of this model is that we let task prices vary across industries, following industry-level production functions. Each product requires different inputs of routine and nonroutine tasks, and routine tasks can be executed by human labour or computer capital. First, I outline the production function that underlies the fundamental dynamics of RBTC, and then I move to more micro-level dynamics by introducing worker-level heterogeneity. The main goal of the model is to outline how technology shapes inequalities between different workers, and how tasks and ICT should play a key role in these inequalities. The implications from the theory form

²Throughout the rest of the paper I will therefore use the terms nonroutine, abstract and cognitive interchangeably. At any time, I use this terminology when referring to tasks that complement software technologies and cannot easily be substituted.

the basis of the later empirical analyses.

2.3.1 Production

The model sets out from the following assumptions. First, computer capital complements nonroutine tasks, yet substitutes routine tasks (Autor et al., 2003). This leads to a production function where routine labour and computer capital are perfect substitutes in executing routine tasks, whereas nonroutine tasks can only be executed by nonroutine labour.

Next, I assume that computer technologies have been differentially adopted across industries in the past decades. This leads to between-industry variation in levels of computers in the current production function. Even though IT is often seen as a General Purpose Technology, and thus that it should affect all industries uniformly, IT adoption is not as pervasive as, e.g. electricity was in the 19th and early 20th century (Jovanovic and Rousseau, 2005). Even though IT shares diffused rapidly in some sectors, other sectors have lagged behind (see also Figure 2.3 in the Data section).

These differences can partially be explained by the suitability of each production process to task replacement or creation. For instance, the routine-replacing nature of computers created more incentives to adopt ICT in routine-intensive industries. Autor et al. (2003) show that industries historically employing large shares of routine labour have seen faster growth in technology adoption than relatively nonroutine-intensive industries. Differences in current-day levels of technology can thus correlate with differences in the routine-intensity of work 30 or 40 years ago. Moreover, the theory of new technology adoption highlights that there always exists random variation in adoption of new technologies, which can depend on uncertainty, sunk costs, average age and size of firms in the industry, export-intensity and possibilities for regional spill-over effects (Haller and Siedschlag, 2011).

In this model, I abstract from analysing why some industries have higher levels of ICT than others. Because I am interested in the worker-level behaviour of sorting into a task-industry combination, I treat industry-level adoption of technology as exogenously given to the worker. They use this information to determine where they will sort into. This thus requires the underlying assumption that workers do not determine the level of technology adoption in an industry, but firms do.³ There is

³There are some other theoretical models that assume that the most skilled workers will adopt technology in their workplace, and that this is the main driver of technology adoption, e.g. Entorf et al. (1999) and Borghans and ter Weel (2007). Therefore, I am assuming that the decision to adopt technology is made by selecting into a job based on the ICT-intensity in the direct environment of that job and comparative

an endogenous entry decision, which we can explain theoretically, and estimate by observing career dynamics of workers switching between industries.

I start from a production function, where each industry j produces a distinct product q_j , which requires a mix of routine and nonroutine tasks. Routine tasks can be executed by computer capital κ and routine labour r , whereas nonroutine tasks can only be produced using nonroutine labour n , due to computer capital being more substitutable for human labour in carrying out routine tasks than nonroutine tasks (Autor et al., 2003). Furthermore, I assume that routine and nonroutine tasks are imperfect substitutes. This results in the following CES production function:

$$q_j = \left[(\kappa_j r_j)^{\frac{\sigma_j-1}{\sigma_j}} + n_j^{\frac{\sigma_j-1}{\sigma_j}} \right]^{\frac{\sigma_j}{\sigma_j-1}} \quad (2.1)$$

Where $\sigma_j = \frac{1}{1-\rho_j}$ is the industry-specific elasticity of substitution and ρ_j the substitution parameter. κ_j is exogenously given and available at no cost, and may therefore also be considered as a technology parameter (Cortes, 2015).

Let $\lambda_{\tau j}$ denote the wage per efficiency unit for task $\tau \in \{r, n\}$ and industry j . Given competitive industries, wages equal the marginal product of labour:

$$\lambda_{rj} = \frac{\delta q_j}{\delta r_j} = \kappa_j^{\frac{\sigma_j-1}{\sigma_j}} r_j^{-\frac{1}{\sigma_j}} \left[(r_j \kappa_j)^{\frac{\sigma_j-1}{\sigma_j}} + n_j^{\frac{\sigma_j-1}{\sigma_j}} \right]^{\frac{1}{\sigma_j-1}} \quad (2.2)$$

$$\lambda_{nj} = \frac{\delta q_j}{\delta n_j} = n_j^{-\frac{1}{\sigma_j}} \left[(r_j \kappa_j)^{\frac{\sigma_j-1}{\sigma_j}} + n_j^{\frac{\sigma_j-1}{\sigma_j}} \right]^{\frac{1}{\sigma_j-1}} \quad (2.3)$$

This allows us to derive relative labour demand for routine tasks, which is given by:

$$\frac{r_j}{n_j} = \kappa_j^{\sigma_j-1} \left(\frac{\lambda_{rj}}{\lambda_{nj}} \right)^{-\sigma_j} \quad (2.4)$$

From the relative demand function, it follows that an exogenous increase in

advantages. Also in terms of timing, this is more sensible. My data set covers a period from 2008 to 2019, which means that ICT is not new anymore, and most adoption has taken place where possible. Workers observe and know this, and sort accordingly.

technology, denoted as an improvement in κ_j , will lead to a decrease in the relative demand for labour performing routine tasks. This is due to the assumption that $\sigma_j < 1$. Rewritten in terms of wage inequality, the main variable of interest:

$$\frac{\lambda_{rj}}{\lambda_{nj}} = \kappa_j^{\frac{\sigma_j-1}{\sigma_j}} \left(\frac{n_j}{r_j} \right)^{\frac{1}{\sigma_j}} \quad (2.5)$$

In the model of Autor et al. (2003) technical change is modelled as an exogenous decrease in the rental rate of capital. In this model, following Cortes (2015), technological change is modelled as an exogenous increase in the stock of computer capital, κ_j . Though the theoretical conclusions do not differ, empirically it is more intuitive to test for complementarity between computers and tasks by measuring the actual computer use at work (i.e. labour-adjusted stocks) rather than its price. I am not modeling an economy-wide increase in technology, but take into account that computer capital might spread unevenly across industries due to i) industry-specific complementarities that arise due to computers being able to perform specific tasks, that might be differentially prevalent in different production processes, and ii) other random variation in the rates of technology adoption in general. Hence the j subscript in the capital indicator.

In the long run, we should expect task prices to equalise across industries, as workers move to industries with higher task prices. But in the short run, there may be considerable scope for task price variation for three reasons. First, task prices are not directly observable, as each occupation is a bundle of different tasks and different industries may offer different bundles. It may take considerable time (and training) for workers with a lower-paying bundle of tasks to switch to an industry and task bundle that has higher rewards. Second, the demand for particular tasks is changing due to technological change and since that process varies across industries, there is scope for industry-specific changes in task prices. Third, supply of specific tasks can also be constrained in the short run, as training of additional workers with particular qualifications takes time.

2.3.2 Worker heterogeneity

Each job comprises a bundle of tasks in which routine or nonroutine tasks are combined. Therefore, the productivity or marginal product of each worker is not either λ_{rj} or λ_{nj} . Furthermore, workers differ in their abilities to perform certain tasks, where some are relatively more productive at performing nonroutine tasks

and vice versa. This motivates a Roy (1951) model of self-selection, similar to Autor and Handel (2013). Because workers are income maximising, they should sort into jobs where they can maximise their returns given their task-specific skills. In terms of terminology, I refer to a job as a specific combination of tasks. An occupation is a pre-defined set of tasks, but as people can change their tasks, there can exist many jobs within each occupation that each have a worker-specific bundle of tasks.

Assume that workers have varying efficiencies or skills in executing different tasks. We can write worker i 's skill endowments as a vector of task efficiencies $\Phi_i = \{\phi_{i1}, \phi_{i2}, \dots, \phi_{iT}\}$. Each element of Φ_i is a strictly positive number measuring the efficiency of worker i at task τ . We can think of Φ_i as a person's stock of human capital, and their efficiency in each task may be a result of human capital investments, innate abilities, or some combination. For the scope of this framework, I only focus on the routine and nonroutine task efficiencies.

From the industry-specific CES production function, we know that each product is produced using a combination of workers and capital. If we assume worker i working on product or industry j is paid their marginal product then their log wage in estimable form should be the following:

$$w_i = \alpha_j + \sum_T \lambda_{\tau j} \phi_{i\tau} + \mu_i \quad (2.6)$$

where $\lambda_{\tau j} > 0 \forall j, \tau$, α_j is an industry specific wage premium and μ_i is a worker-specific error term. We can simplify (2.6) to accommodate for the fact that there are only two tasks in the economy by writing

$$w_i = \alpha_j + \lambda_{rj} \phi_{ir} + \lambda_{nj} \phi_{in} + \mu_i \quad (2.7)$$

Note that the wage of each worker is not only dependent on their specific efficiencies in the task, but also on the production process of the industry in which they are employed. From (2.3) and (2.2) it follows that $\lambda_{\tau j}$ includes the level of ICT κ_j that is present in the production process these workers operate in, as well as the general (non)routine task-intensity of that industry. Therefore, the task prices - and thus the individual's wage - will depend on the level of routine and nonroutine tasks that are used in the specific industry, as well as the level of computer capital.

This economy is therefore characterised by self-selection of workers into industries, based on comparative advantage. In equilibrium, the marginal worker in industry j is indifferent between that industry and the next best alternative. The equilibrium of the model ensures that workers are employed in the industry that has

the highest reward to their bundle of tasks, given their task efficiencies $\phi_{i\tau}$ and task returns λ_{τ_j} .

Note that at this point, this paper differs from almost all other RBTC-related literature in that it does not observe sorting at the occupation-level, but the industry-level. Because this paper is able to rely on individual-level data, I do not need to observe the occupation, and thus assume that all workers in an occupation perform the same set of tasks. Therefore, we can see how people use their tasks in different production processes in order to establish a direct link between individual tasks and the task inputs in a production process, and thus also more precisely measure the complementarity (substitutability) between technology - in the shape of ICT - and nonroutine (routine) tasks.

2.3.3 Empirical implementation

Cross-sectional implications

Identifying the market average of the returns to skill in the presence of self-selection is not empirically straightforward. As outlined in the model, I expect that workers non-randomly sort into industries. Thus, a regression of log wages on workers job tasks will not generally recover the actual average returns or prices to these tasks. Workers with high productivity for certain tasks will sort toward jobs that highly reward those tasks, due to self-selection on comparative advantage. Therefore, a Mincer (1974) type regression using OLS estimation on worker-level nonroutine N_i and routine R_i tasks based on the form,

$$\ln w_i = \alpha + \beta_N N_i + \beta_R R_i + e_i \quad (2.8)$$

would lead to biased estimates, because the coefficients for each task might include self-selection based on task-efficiencies. It is very likely that β_N and β_R would not solely capture the “task-return” across the economy, because a) it varies between industries, and b) there is non-random sorting of workers into tasks based on comparative advantage. Though useful for descriptive analyses, an estimation of the Mincer equation in the above form would not illuminate the full story on how changing technology impacts outcomes for workers with different comparative advantages across tasks.

Therefore, there are a number of assumptions that should be dealt with, as discussed in detail by Autor and Handel (2013). First, workers can only sort into industries where their returns to the bundle of tasks is larger than the returns in

another industry. For the same bundle of tasks, workers will always sort into the industry that pays the higher reward. Second, tasks must negatively covary within the set of industries that have positive employment. Workers will self-select into jobs that differentially reward tasks that they have comparative advantage in. Third, if the correlation between workers in terms of skill is not too high, workers will self-select into jobs that offer high returns to the tasks in which they are particularly efficient - meaning that self-selection takes the form of comparative advantage.

In practice, it is difficult to test for sorting into jobs based on complementarities between workers' task efficiencies and the task demanded for a certain job (Eeckhout and Kircher, 2011). To guide the empirical identification of sorting patterns according to comparative advantage, one of the possibilities the data provides is to test for the difference in returns to tasks between industries that vary in the intensity of the task used. In other words, if tasks returns are higher in industries that use that task intensively, this would indicate positive sorting on comparative advantages in (non)routine task efficiencies.

To partly overcome the issue of self-selection, I apply a similar method as Autor and Handel (2013). Sorting on comparative advantage would relate to the empirical observation of nonzero covariances between industry-level task returns and the task endowments of workers who self-select into these industries (Autor and Handel, 2013). To recover these covariances, we can estimate an augmented version of Equation (2.8) where we interact industry-level task means $\bar{R}_j$ and $\bar{N}_j$ with worker-level task inputs:⁴

$$\ln w_{ij} = \alpha + \beta_N N_i + \beta_R R_i + \delta_N \bar{N}_j + \delta_R \bar{R}_j + \gamma_N N_i \times \bar{N}_j + \gamma_R R_i \times \bar{R}_j + e_{ij} \quad (2.9)$$

Two cases emerge that make different predictions on the signs of the interaction terms γ_N, γ_R in this equation. The first case is one of *comparative advantage*, where workers positively self-select into each bundle of tasks. In terms of the data presented in this paper, this would be when a worker who performs relatively more routine-intensive tasks than the average worker, earns more when he works in an occupation or industry where the routine-intensity of work is high. This can occur when the correlation between worker abilities across tasks is sufficiently low: workers who are able to perform abstract tasks are not necessarily also productive in routine tasks, and vice versa. In other words, routine and abstract tasks are not on the same skill spectrum, where high-skilled workers perform abstract tasks and low-skilled

⁴Note that Autor and Handel (2013) use this same equation, but group workers at the occupation level. In analysis I will estimate this regression with both occupation averages and industry averages in order to validate my results and compare them with their paper.

workers routine tasks.⁵ This implies that there will be a positive covariance between industry-level task returns and the worker-specific task endowments. Task returns should be higher in industries that differentially use these tasks. Formally, this implies that $\gamma_N, \gamma_R > 0$.

The other case is one where the distribution of skills across the population of workers is characterised by *absolute advantage* - that is, workers who excel at task 1 also excel at task 2 - then positive self-selection on task 1 into industry j must imply negative self-selection on task 2 into industry j' and vice versa. For instance, if nonroutine tasks are simply more complex tasks than routine tasks, and they require the same skill but at a higher complexity, we could observe positive self-selection of workers in nonroutine tasks and negative self-selection in routine tasks. Hence, if all the able workers have self-selected into nonroutine tasks, then the more routine-intensive industries will have workers who are relatively unproductive at the tasks that are differentially rewarded in that industry, being routine tasks. In the regression coefficient, this would be picked up in the fact that γ_N is positive, but that γ_R might be negative.

The main issue that I try to solve by estimating Equation (2.9) is to disentangle task rewards from sorting patterns. In the Mincer equation (2.8) alone, I estimate the β_T coefficients, that could be interpreted as rough task prices. However, in light of a Roy model of sorting according to comparative advantage, the β_s are biased and might pick up the fact that able workers sort into nonroutine tasks. Therefore, a high coefficient on nonroutine-intensity might be the result of non-random sorting patterns. To determine whether sorting is happening - on both the occupation and the industry level - I estimate equation (2.9). The returns to individual tasks that we find here should not reflect the sorting-effect, which implies that they are more reflective of the 'true' relation between routine work and wages. If γ_R and γ_N are positive, then we have the situation of comparative advantage: workers self-select into industries where they can maximise their return, and we thus observe that workers that perform routine tasks also have higher task returns in industries where there are more routine tasks executed.

Estimating a regression in the spirit of Equation (2.9) therefore allows us to partly disentangle $\phi_{i\tau}$ and $\lambda_{\tau j}$. However, we can pick up more from the task price element $\lambda_{\tau j}$ if we specifically also regress the technology inputs from (2.3) and (2.2). In order

⁵Specifically, the sufficient condition for this to occur is that the correlation ρ between abilities in tasks is such that $\rho < \min\left(\frac{\lambda_n}{\lambda_r}, \frac{\lambda_r}{\lambda_n}\right)$ or at least that $\rho \leq 0$ which implies that worker abilities in tasks n and r are either uncorrelated or negatively correlated.

to observe the complementarities between the amount of computer capital K_j ⁶ that is adopted in an industry and the tasks of the workers in that industry, we can estimate:

$$\ln w_{ij} = \alpha + \theta_K K_j + \sum_{T \in (N,R)} [\beta_T T_i + \theta_T K_j \times T_i] + e_{ij} \quad (2.10)$$

where T_i captures (either nonroutine or routine) worker-level tasks. This is a novel way to use the Mincer equation in an RBTC setting. Normally workers would have sorted into occupations that are assumed to remain constant in tasks. In this paper, I highlight that tasks can vary within occupations, and workers will provide their tasks to the production process (i.e. industry) where they will maximise their returns. If this is the case, we should see that in high ICT-intensive industries, the return to abstract work should be positive and higher than in non-ICT-intensive industries. Following the same, but inverse, line of reasoning, routine work should have lower task returns in ICT-intensive industries. Also, ICT capital in general should be beneficial to all wages, because κ is positively related to wages, in the equation of $\lambda_{\tau j}$ (Eq. (2.2) and (2.3)).

Longitudinal implications

The above-described equations all account for cross-sectional data. The main advantage of the data used in this paper is that we can also observe worker careers in terms of job changes and wage growth. This allows for a more detailed analysis of the sorting patterns and wage inequality in the context of the differential spread of technological progress. Even though the task data remain cross-sectional (as people are only surveyed once), I can retrieve these individuals in administrative data files where I have longitudinal information about a) their wages, b) their employed industry, c) the industry-average task-intensity and d) the industry-average ICT-intensity in a period of 1 to 8 years after the survey (see Section 2.4 for a further discussion of the data sources).

Following the Roy model of self-selection, we should expect to see that workers will switch jobs in order to maximise their earnings over their career. Think of a simple two-period setting, in which a worker has task efficiency $\phi_{i\tau}$ and searches for the highest possible $\lambda_{j\tau}$. When making the decision to stay in industry j or switch to industry j' , there exist ‘push’ and ‘pull’ factors. The former are those characteristics of a worker combined with a certain work environment, that make a worker more likely to be searching for a switch. For instance, highly educated nonroutine workers

⁶Note that the model notation κ_j changes to K_j in the estimation, because computer use is a proxy for κ_j .

should be more likely to switch, and relatively well-paid routine workers should be more likely to stay given that they can probably not find a pay increase elsewhere. Following the empirical findings by Cortes (2015), we can expect that there are certain characteristics that make a person more likely to make a switch.⁷ Exit probabilities increase for the relatively more and the relatively less able workers, meaning that workers are most likely to flow out an industry when they are in either of the tails of the ability distribution.

The ‘pull’ factor is the setting where a worker will compare the expected returns across industries. In this case, adopting the notation from Equation (2.7), a worker will switch to a certain industry j' if

$$(\alpha_j - \alpha_{j'}) + (\lambda_{rj} - \lambda_{rj'})\phi_{ir} + (\lambda_{nj} - \lambda_{nj'})\phi_{in} > 0 \quad (2.11)$$

If workers will sort according to comparative advantage, we should expect that routine workers sort to more routine-intensive industries, and vice versa for nonroutine workers.

Besides switching patterns, we can also observe wage growth over the period after the survey. However, as I can only observe the worker during the survey at time t_0 , I can only include T_i and no time-varying component of task endowments. It should therefore be seen as an indicator of path dependency: if workers have sorted into a routine or abstract-intensive job, this will have an impact on the growth factor of their wages in the coming period. The function that I will be estimating takes the following growth form, where the superscript L denotes a longitudinal coefficient

$$\ln w_{ijt+x} - \ln w_{ijt} = \alpha + \beta_R^L R_i + \beta_N^L N_i + \varepsilon_{ijt} \quad (2.12)$$

Lastly, we should expect that a worker’s wage growth is dependent on both their starting industry, as well as the switches they make, and specifically on the direction of the switch, as in Cortes (2015). For the interaction between the industry where the worker resides at the time of the survey, and how this potentially relates to future wage growth, I can estimate the following separately for $I \in \{K, R, N\}$:

$$\ln w_{ijt+x} - \ln w_{ijt} = \alpha + \sum_{T \in (n,r)} \left(\beta_T^L T_i + \delta_T^L \bar{I}_j + \gamma_T^L \bar{I}_j \times T_i \right) + \varepsilon_{ijt} \quad (2.13)$$

However, it is likely that a person switches industries in the period following the survey. For the switching patterns we are interested in wage growth in the period

⁷Even though Cortes (2015) discusses switching occupations, the dynamics should be similar.

of $t + x$ years, conditional on (not) having switched career track. For this equation we introduce a dummy $S_{ij'jt+x}$ that captures whether a person has switched from industry j to industry j' in the period $t + x - t_0$. The dummy can take three values: 1 if a person remains in the same industry, 2 if a person moves upward in ICT-intensity of the industry, and 3 if a person moves downward. I recreate this dummy for both the ICT-intensity, as well as the average routine and nonroutine-intensity of the industry.

$$S_{ij'jt+x} = \begin{cases} 1, & \text{if } j_{t+x} = j_{t_0} \\ 2, & \text{if } K_{jt+x} > K_{j_{t_0}} \\ 3, & \text{if } K_{jt+x} < K_{j_{t_0}} \end{cases} \quad (2.14)$$

We end up with the following function (where $S_{ij'jt+x}$ is simplified to S for ease of notation):

$$\ln w_{ijt+x} - \ln w_{ijt} = \alpha + \beta_R R_i + \beta_N N_i + \zeta_S S + \zeta_R S \times R_i + \zeta_N S \times N_i + X_{it} \beta_3 + Z_{ijt} \beta_4 + \varepsilon_{ijt} \quad (2.15)$$

2.3.4 Empirical predictions

This framework produces a number of empirical implications. First, it predicts that nonroutine tasks are associated with higher wages than routine tasks. This is in line with findings using both occupation-level task data as well as individual-level task data. The novel aspect of my approach is the fact that we can discriminate workers within occupations in terms of tasks, and due to the large sample size can also observe between-industry differences in tasks. Therefore, we can make predictions about the role of ICT in the formation of wage inequality. Specifically, given the complementarity between computer capital and nonroutine tasks on the industry level, we should expect that inequality between routine and nonroutine workers should be higher in the presence of more ICT.

Second, we can observe switching patterns and analyse whether workers move according to their revealed comparative advantage. This allows us to see if relatively more (non-)routine workers are more likely to switch to i) nonroutine-intensive industries, ii) routine-intensive industries or iii) ICT-intensive industries. We should expect to see that nonroutine workers should move to relatively more nonroutine industries, if they switch, and routine workers to routine industries. Furthermore, given RBTC, we should also expect that relatively more nonroutine workers should move to industries where their tasks are more complementary to, and thus to more

Table 2.1: Empirical Predictions for i) worker-level measures, ii) industry-level averages of tasks and ICT and iii) interaction terms.

Estimation	Eq.	Worker-level				Industry-level				Interaction terms		
<i>Cross-sectional</i>												
Baseline	(2.8)	β_N	+	β_R	-							
Comparative advantage	(2.9)	β_N	+	β_R	-	δ_N	+	δ_R	-	γ_N	+	γ_R +
ICT interactions	(2.10)	β_N	+	β_R	-	θ_K	+			θ_N	+	γ_R -
<i>Longitudinal</i>												
Baseline	(2.12)	β_N^L	+	β_R^L	-							
Industry interactions: $\bar{N}_j$	(2.13)	β_N^L	+	β_R^L	-	δ_N^L	+			γ_N^L	+	γ_R^L -
Industry interactions: $\bar{R}_j$	(2.13)	β_N^L	+	β_R^L	-			δ_R^L	-	γ_N^L	-	γ_R^L +
Industry interactions: $\bar{K}_j$	(2.13)	β_N^L	+	β_R^L	-	δ_K^L	+			γ_N^L	+	γ_R^L -

ICT-intensive industries. The converse should hold for routine workers.

Lastly, conditional on switching to either more or less routine, nonroutine or ICT-intensive industries, we should see that wage growth acts accordingly. Routine workers who move to an even more routine-intensive industries are likely to also have lower wage growth than nonroutine workers, due to the lower task price of routine work. However, if they switch to a more nonroutine industry this most likely also relates to a change in tasks, and should thus relate to an increase. Most likely, here we should see that human capital variables influence the probability of the direction of the switch as well as being a good predictor for later wage growth, conditional on switching.

If there is indeed routine biased technological change we should expect the following outcomes, also described in Table 2.1. Nonroutine tasks should have positive returns, and routine tasks negative.

If there is indeed matching of workers and industries based on comparative advantage in the two routine and nonroutine task efficiencies, we should expect the following. The covariance between individual tasks and industry-mean tasks should be positive: task returns are higher in industries where that task is used more. Moreover, conditional on switching, nonroutine workers should move to nonroutine industries, and vice versa. If neither of these predictions hold, than any predictions on RBTC should be treated with caution, as the effects might thus capture general ability rather than task efficiencies.

Table 2.2: Empirical Predictions on Switching Patterns and Wage Growth Conditional on Switching.

	to n_j^{HI}	to r_j^{HI}	to K_j^{HI}	Theoretical channel
Probability of switching				
n_{ij}	+	-	+	Unclear sign: despite high ϕ_{ir} , $\lambda_{nj} > \lambda_{rj}$ so $\lambda_{nj'}\phi_{nj'} + \lambda_{rj'}\phi_{rj'} \gtrless \lambda_{nj}\phi_{nj} + \lambda_{rj}\phi_{rj}$
r_{ij}	+/-	+/-	+/-	
Pull factors				
$\bar{n}_j - n_{ij}$	+	-	+	High difference should make workers more likely to switch, but not to more routine-intensive industry: complementarity between K and n
$\bar{r}_j - r_{ij}$	-	+	-	Substitutability between K and r , so this should not be a factor
$\bar{educ}_j - educ_{ij}$	+	-	+	Workers want to match their education to the average education level of the group. If this difference is larger, they should be less inclined to move.
Push factors				
$n_{ij} - \bar{n}_j$	+	-	+	Relatively more nonroutine workers within an industry should be more likely to switch a more nonroutine-intensive industry, and complementarity with ICT
$r_{ij} - \bar{r}_j$	-	+	-	Similar to the above, but also substitutability with ICT
$educ_{ij} - \bar{educ}_j$	+	-	+	High ability workers should move to more nonroutine, ICT-intensive industries, and conversely for less able workers
Conditional wage growth				
n_j	+	-	+	Unclear sign: despite high ϕ_{ir} , $\lambda_{nj} > \lambda_{rj}$ so $\lambda_{nj'}\phi_{nj'} + \lambda_{rj'}\phi_{rj'} \gtrless \lambda_{nj}\phi_{nj} + \lambda_{rj}\phi_{rj}$
r_j	+/-	+/-	+/-	

2.4 Data

This paper takes advantage of data from the Netherlands Working Conditions Survey (NEA). The NEA is an annual survey that has been running since 2003. It is nationally-representative of workers aged between 15 and 75. It polls around 23,000 to 30,000 workers each year, and surveys are sent to people's home addresses, and can be either filled out on-line or using a paper version. The task measurements in this paper are based on several questions from the working conditions section, which are available from 2008 onward, with the exception of 2013 and 2015. One benefit from this survey is that it asks workers both directly for their job tasks and for their occupation, creating data suitable for understanding within-occupation variation in tasks. Moreover, NEA contains demographic variables such as education level, age, gender, and cultural background, and company and industry data, which are not self-reported but originate from census data. Descriptive statistics are presented in

Table 2.4.⁸

The main benefit of using microdata from Statistics Netherlands is that all citizens have an anonymised personal identifier, based on their social security number, which allows for matching across all sorts of data files. As NEA is incorporated in this microdata system, it is possible to retrieve data on income and employment from census records (called POLISBUS), which is provided by employers to the tax agency. Matching survey data to census data provides more accurate income information compared with self-reported surveys (Abowd and Stinson, 2013; Meyer and Mittag, 2019a,b), but it also allows for following workers over time. At the time of writing, I can make use of the POLISBUS for up to and including 2019. Because of COVID-19, 2019 will mark a natural ending point and even though the data is available, I have not included 2020. The POLISBUS data consists of monthly information on hours worked, gross and net wages, type of contract⁹ and unemployment spells.¹⁰ Matching occurs based on both the personal identifier as well as a unique job identifier, which is based on the employer-employee combination. This extra matching requirement is necessary, as some workers have several jobs at time of the survey. In NEA, respondents are asked to fill out the survey for the job they spend the most hours in. As POLISBUS contains working hours data, I observe which job has the most contracted hours in the period the survey was open and match the survey to that job.

This paper uses seven pooled cross-sections from NEA: five waves from 2008 to 2012, plus 2014 and 2016. From 2008 to 2013, each cross-section contains between 20,000 to 25,000 respondents. After 2014, the size of the survey has increased and in these the sample size grows to a size of 30,000 to 35,000 workers. I remove respondents that are younger than 25 or older than 65, those that cannot be matched to census data with certainty (because they either switched jobs during the survey period, did not have an official taxable job, or had two jobs with the same hours), and those with missing data on task measures or demographic variables. The final sample covers 154,398 workers. Only survey wave 2014 and 2016 contain occupation

⁸The data are made nationally-representative due to the inclusion of sampling weights, provided by Statistics Netherlands. The weight coefficients are constructed using post-stratification. Stratifications used are: 1. gender x age cohort x migration status, 2. industry, 3. region x urban and 4. gender x age cohort x level of education. The sampling weights are used in the final estimations. See Hooftman, Mars, Janssen, Vroome, Ramaekers, and Bossche (2018) for further explanations.

⁹This can be either tenured or temporary, where a tenured contract applies to workers with a contract for an indefinite period, plus interns, directors/major shareholders ('directeur-groootaandeelhouder or dga in Dutch), and people employed under the Sheltered Employment Act (wsw in Dutch). Temporary contracts apply to temporary employees, sub-contracted or on-call employees (uitzendkracht and oproepkracht in Dutch, respectively)

¹⁰I create the unemployment manually, by filling up the empty months in between contracts with a dummy when there is missing data. The hourly wage is obtained by dividing the gross monthly wage by contracted hours.

Table 2.3: NEA task measures by subgroups

		Total	Gender		Education		
			Male	Female	Low	Middle	High
Abstract-intensity							
Control over methods	Often	67.6	73.1	61.8	57.2	62.9	76.6
	Sometimes	22.7	19.8	25.7	26.1	25.7	18.3
	No	9.8	7.1	12.6	16.7	11.5	5.1
Problem solving	Often	74.2	79.1	69.0	56.6	70.4	85.5
	Sometimes	22.8	18.5	27.3	35.8	26.6	13.4
	No	3.0	2.4	3.7	7.6	3.0	1.1
Learning	Often	50.6	51.6	49.5	29.6	45.0	65.2
	Sometimes	42.9	42.9	43.0	55.2	48.6	32.1
	No	6.5	5.5	7.5	15.3	6.4	2.7
Routine-intensity							
Control over speed	Often	62.1	67.0	56.9	57.5	59.8	66.3
	Sometimes	25.1	23.7	26.6	26.5	26.6	23.1
	No	12.8	9.3	16.5	16.0	13.6	10.6
Repetitive movements	Often	31.1	30.5	31.7	48.4	35.7	19.1
	Sometimes	20.4	21.6	19.2	23.1	23.0	16.8
	No	48.5	48.0	49.0	28.5	41.3	64.1
Control over sequence	Often	69.4	72.3	66.4	58.2	66.0	77.6
	Sometimes	20.2	19.0	21.5	24.0	22.6	16.2
	No	10.4	8.7	12.2	17.8	11.4	6.2
Nonroutine score (standardised)		0.01 (1.00)	0.13 (.94)	-0.13 (1.03)	-0.52 (1.13)	-0.11 (1.00)	0.35 (.78)
Routine score (standardised)		0.00 (1.00)	-0.10 (.95)	0.10 (1.04)	0.32 (1.06)	0.10 (1.02)	-0.24 (.89)

Note: Note: Data from NEA. The table shows the percentage of respondents within each group. Education groups are: low (ISCED 0-2), middle (ISCED 3-4) and high (ISCED 5 and up). (Non)routin scores are calculated using polychoric PCA, with sampling weights. See Appendix 2.A for exact phrasing of survey questions.

data using ISCO classifications, in contrast to the first cohorts that use a classification of occupations that does not allow for international comparison.

2.4.1 Task and ICT Measures

The main rationale behind studying tasks in relation to technological progress, is that computer capital can substitute for workers that perform limited and well-defined set of tasks. At the same time, it complements workers in carrying out problem-solving tasks, as the information they can solve problems with is more readily accessible Autor et al. (2003). Following Acemoglu and Autor (2011), routine tasks are ‘procedural, rule-based activities to which computers are well-suited’, whereas nonroutine tasks are ‘activities that require problem-solving, intuition, persuasion

and creativity’.

Note that this paper does not distinguish between routine manual and routine cognitive, nor between abstract cognitive and abstract interactive. Rather, I use routine and abstract as two broader terms that refer to either substitutable by machines, or complementary¹¹. I define routine tasks as those tasks that can be described using rules and explicit procedures (Autor et al., 2003). The worker does not make these decisions themselves but rather follows instructions. The construction of an index for abstract work is based on problem-solving, decision-making and learning. These characteristics reflect need for workers to adapt to changing circumstances, rather than following guidelines. It thus captures nonroutine-intensity. Throughout, I will use the terms nonroutine and abstract interchangeably.

I rely on three items from NEA to construct a nonroutine task measure: (1) Can you decide how to execute your work?; (2) In your work, are you required to think of solutions to do certain things?; and (3) Does your job require learning new things? This reflects that nonroutine cognitive tasks require problem-solving and creative skills, where workers are autonomous in deciding how to do tasks, rather than following explicit rules. To measure routine-intensity, I use the following questions: (1) Are you doing work where you are required to make repetitive movements? and (2) Can you decide the sequence of your tasks? and (3) Can you decide the speed of your tasks? These inputs reflect the explicit procedure characteristic of routine tasks (Autor, Levy, Murnane, 2003).

The questions can be answered by ‘no’ (response category 3), ‘yes, sometimes’ (2) and ‘yes, often’ (1). Table 2.3 shows the distribution of answers to these questions across subgroups. Note that I observe the routine-intensity of one’s entire job, rather than time spent on or the importance of specific tasks, such as in Stinebrickner et al. (2018), Autor and Handel (2013) and O*NET questions. It is a reflection of the perception of workers in terms of how routine and nonroutine their tasks are. Naturally, this may be clouded by personal differences in the interpretation of questions. However, I have no reason to suspect a systematic pattern in those differences.

The data from NEA give an indication of the routine-intensity of jobs in the Netherlands. Table 2.3 shows that higher educated workers have more opportunities to learn on the job than middle educated and lower educated workers. Likewise, they are also more likely to have autonomy over the way they perform their tasks

¹¹I refrain from nonroutine manual and service tasks altogether, as these are tasks not directly impacted by the increase of technology. Rather, workers executing these tasks are affected as part of equilibrium effects where demand for low-skill services increase because of increased wages for high-paying jobs, see e.g. Cortes (2015) and Gregory, Salomons, and Zierahn (2021)

and can more often be found solving problems. Lower educated workers are to a larger degree responsible for performing tasks that require little abstract thinking. At the same time, their job descriptions contain more programmable tasks: they have lower control over the speed and sequence of their tasks and are more likely to do repetitive movements than higher educated workers. Furthermore, men are required to learn more, solve more problems and have more control over their work methods than females have. Similarly, they perform fewer routine-intensive tasks.

Based on these variables, I construct a standardised¹² index, using Principal Component Analysis. Standard PCA relies on Pearson correlations, which assumes that all variables are normally distributed. This poses a problem for the variables proposed here, which have a discrete scale. Also, I would need to assume that the distance between all the answers are equal, and thus that the step from 'no' and 'yes, sometimes' and 'yes, sometimes' and 'yes, often' is the same (Bond and Lang, 2018). To solve this, I create an index using PCA based on polychoric correlation, which assumes the variables are ordered measurements of an underlying continuum (Olsson, 1979a). This makes it better suited for creating an index using categorical variables. The routine index captures 62% of variation in the underlying variables, and for the abstract measure it captures 59%, as can also be seen in Table 2.A.1.¹³

To show that these measures capture the same underlying notion of routine and abstract-intensity as is common in the literature, I compare my measures to the routine and abstract-intensity scores index using O*NET data. Table 2.A.1 provides the routine and abstract scores based on occupational dictionaries and the occupational averages of the standardised scores from NEA and EWCS data, for 2 digit occupations. Occupational averages of the standardised indices are comparable with O*NET based measures, as used by Autor and Dorn (2013), Goos et al. (2014) and Cortes (2015). The correlation coefficient between the routine score from NEA and the Goos and Salomons (2009) measure is 72% , and 84% for the nonroutine score.

However, the PCA-based index only provides an indication of the relative abstract and routine-intensity of one's job, and not the absolute routine content of their job. To get an idea on 'how' routine-intensive the sample population is, I also create an alternative typology of routine-intensity, where each worker is classified as being either routine, abstract or a 'mix'. The goal of creating this typology is to put the PCA into perspective, as the standardised measures of routine-intensity might give

¹²Standardisation occurs over the entire population, by survey year

¹³I weighted the polychoric PCA using sampling weights provided by CBS, in order to create indices that are representative for the entire population rather than just the survey sample.

Figure 2.2: Comparison of routine-intensity in occupations, in relative terms using PCA and in absolute terms using the typology

Source: Author's calculations using NEA data.

a skewed image of the absolute routine-intensity of the Netherlands. In order to classify one as being routine, I require that they answer at least one question with 'yes, often' and no other questions within the same component with a 'no'. In this way, I conclude that the main tasks of the worker are routine and thus that he is in a routine job. I also construct this measure for the nonroutine-intensity, where I impose the requirement that a worker cannot be in a job that is both routine and abstract. Workers that do not fit this description are put in the mix-category.

To highlight this point, Figure 2.2 shows the standardised indices and the typology. Observing the left panel, the expected pattern emerges: managers are more abstract-oriented and not routine-intensive, whereas workers in elementary occupations or plant and machine operators do not have abstract, but very routine-intensive jobs. This pattern is largely consistent over the increase in complexity that underlies the ISCO classification. Conversely, the right panel shows that even though the PCA shows large differences between occupations, this does not necessarily imply that workers in lower-level occupations are routine-intensive per se. The largest share of Dutch workers fall into the category of Mix or Abstract type workers, rather than that of Routine workers, regardless of their occupation. The Netherlands does not seem to be a very routine-intensive country, but rather abstract-oriented. In other words, many routine tasks have likely already been replaced by machines, but a

stratum of these tasks remain to be executed by human labour.

These shares of routine workers seem small, but a similar analysis using occupation-level data would not lead to drastically different conclusions. Table 2.A.2 in the Appendix shows the share of Dutch workers in a routine, abstract or mix job and the share of workers in level 1 to 4 ISCO occupations, ranked by complexity and routine-intensity. Level 1, occupations with routine tasks and little complexity, comprise only 5% of the labour force. Similarly, routine workers from the typology make up for 5.4%. Interestingly, these do not seem to be exactly the same workers. Roughly 5% (3/5.4) of routine workers can be found in ISCO level 1 occupations, but more than 50% (2.8/5.4) of them are located in a level 2 occupation, which is larger than the global average of 38%. All in all, both occupation as well as individual-level data show that the Netherlands is a predominantly abstract-oriented economy.

For the ICT indicator, I use a question on computer hours. Workers are asked how many hours a day they spend working on a computer (which could also be a smartphone, tablet or laptop, and not necessarily a desktop). I use the industry-level average of these hours, and standardise this to create an ICT indicator with mean 0 and standard deviation 1. The computer use data shows that there are significant differences between industries (see Figure 2.3).¹⁴ These variations can be exploited in the analysis of wage differences conditional on ICT adoption in the industry as well as differential probabilities of switching to a more or less ICT-intensive industry.

I choose to use industry-level data rather than firm-level or solely individual-level data on computer use for three reasons. First, there is substantial and sufficient variation between industries (as can also be seen in Figure 2.1). Second, the only firm-level microdata on ICT available consists of either investment flows or the value of the ICT stock in a company. However, this is not informative about the actual use and thus adoption of technologies in the workplace. A company might have invested in expensive computers, and still uses its possibilities to the same extent as a firm who invested in second-hand computers. Third, I can construct the industry-data on computer use using individual-level survey data on computer use from the same survey as the tasks. Therefore, I can also directly include a control for workers who use computers above industry-average, and are thus successful adopters. Aggregating computer use on the firm-level would rely on too narrow bins in terms of observations and would thus require the rather restrictive assumption that each worker we observe in the survey is representative for his or her entire firm.

¹⁴In the analyses I use measures for ICT-intensity on a year-industry level.

Table 2.4: Descriptive Statistics

		Total	Gender		Education		
			Male	Female	Low	Middle	High
Demographics							
Age	Age	44.99	46.05	43.88	49.16	45.03	43.11
	<i>sd</i>	(10.65)	(10.61)	(10.58)	(9.7)	(10.34)	(10.81)
Education	Low	.18	.19	.17	1	0	0
	Middle	.40	.40	.40	0	1	0
	High	.42	.41	.43	0	0	1
Native	Dutch	.86	.87	.86	.85	.87	.86
	Non-Dutch (western)	.08	.07	.08	.07	.07	.09
	Non-Dutch (non-western)	.06	.06	.06	.08	.06	.05
Job characteristics							
Task scales							
	Abstract	.01	.13	-.13	-.52	-.11	.35
	<i>sd</i>	(1.00)	(.94)	(1.03)	(1.13)	(.99)	(.78)
	Routine	.00	-.10	.10	.32	.10	-.24
	<i>sd</i>	(1.00)	(.95)	(1.04)	(1.06)	(1.02)	(.89)
Log hourly wage		2.96	3.06	2.86	2.71	2.86	3.18
	<i>sd</i>	(.44)	(.46)	(.39)	(.34)	(.37)	(.44)
Hours		133.23	153.25	112.25	128.24	131.6	136.98
	<i>sd</i>	(41.4)	(31.68)	(39.98)	(46.49)	(42.43)	(37.52)
Non-tenured contract (months in t_0)		.39	.34	.46	.61	.45	.24
	<i>sd</i>	(2.03)	(1.87)	(2.18)	(2.52)	(2.17)	(1.58)
Unemployed months in t_0		.33	.33	.33	.42	.32	.31
	<i>sd</i>	(1.45)	(1.46)	(1.45)	(1.62)	(1.44)	(1.38)
Career progression							
Wage growth							
	$\ln w_{t+2} - \ln w_{t_0}$	.03	.03	.03	.01	.02	.04
	<i>sd</i>	(.23)	(.24)	(.22)	(.21)	(.22)	(.24)
	$\ln w_{t+5} - \ln w_{t_0}$	.04	.04	.05	.01	.03	.07
	<i>sd</i>	(.28)	(.30)	(.26)	(.26)	(.26)	(.30)
	$\ln w_{t+8} - \ln w_{t_0}$	.06	.05	.07	.02	.04	.10
	<i>sd</i>	(.32)	(.35)	(.29)	(.29)	(.29)	(.35)
Switching to higher							
	$\bar{R}_j > R_i$	.11	.12	.10	.10	.11	.11
		(.31)	(.32)	(.30)	(.30)	(.31)	(.31)
	$\bar{N}_j > N_i$	.54	.65	.43	.38	.47	.67
		(.50)	(.48)	(.49)	(.49)	(.50)	(.47)
	$\bar{K}_{j'} > K_j$	.11	.12	.10	.10	.10	.11
		(.31)	(.32)	(.29)	(.30)	(.31)	(.31)
Switching to lower							
	$\bar{R}_j < R_i$	.10	.11	.10	.09	.10	.11
		(.30)	(.31)	(.30)	(.29)	(.30)	(.31)
	$\bar{N}_j < N_i$	.11	.12	.10	.10	.11	.11
		(.31)	(.32)	(.30)	(.30)	(.31)	(.31)
	$\bar{K}_{j'} < K_j$	.11	.11	.10	.09	.10	.11
		(.31)	(.31)	(.30)	(.29)	(.31)	(.32)
Observations		155,781	79,718	76,063	28,603	62,143	65,035

Note: Data from NEA. Standard deviations (sd) in brackets. Switching is a dummy equal to 1 if industry in $t + 5$ is different from industry in t_0 . Education groups are: low (ISCED 0-2), middle (ISCED 3-4) and high (ISCED 5 and up). Native is classified as Dutch if person is Native, (non-)western if a person has a migration background from either a western or non-western country.

Figure 2.3: ICT-intensity by one digit industries

Source: Author's calculations using non-public NEA data from Statistics Netherlands.

Note: ICT-intensity measured as average hours of daily computer use, which is standardised over entire population. Sector specifications (low skill services, goods, and high skill services) used following Bárány and Siegel (2020) to visualise differences between industries in the same sector.

2.5 Results

2.5.1 Baseline regressions

For the baseline descriptive cross-sectional equation, I estimate an OLS regression for log hourly (monthly) wage $\ln w_{ij}$ of individual i in industry j as follows:

$$\ln w_{ij} = \beta_0 + \beta_1 R_{ij} + \beta_2 N_{ij} + X_i \beta_3 + Z_i \beta_4 + \varepsilon_i \quad (2.16)$$

where R_{ij} and N_{ij} are the standardised routine and abstract task scales. X_i is a vector of demographic controls (education, age, age², gender and migration background) and Z_i is a vector of job-related controls (occupation and tenure), ε_i is the worker-specific error-term. Each regression includes a dummy of the year of the survey. Including basic socio-demographic variables partly accounts for selection on observables.

The results are presented in Table 2.5. All standard errors are robust and clustered

Table 2.5: OLS Regressions of Log Hourly Wages on Task Intensities, Demographic Variables and Occupation Dummies

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Routine	-.031*** (.0014)	-.026*** (.0013)	-.019*** (.0012)	-.012*** (.0012)	-.0019 (.0017)	-.0026 (.0016)	.0050 (.0034)	-.014*** (.0030)
Nonroutine	.14*** (.0014)	.10*** (.0013)	.076*** (.0013)	.066*** (.0012)	.048*** (.0017)	.041*** (.0016)	.14*** (.0040)	.22*** (.0028)
Middle educ			.080*** (.0027)	.11*** (.0027)	.070*** (.0040)	.062*** (.0040)	.14*** (.0072)	.098*** (.0057)
High educ			.33*** (.0033)	.39*** (.0032)	.24*** (.0052)	.22*** (.0052)	.43*** (.0096)	.28*** (.0074)
Female				-.15*** (.0022)	-.12*** (.0034)	-.10*** (.0035)	-.41*** (.0097)	-.42*** (.0052)
Age				.053*** (.00082)	.057*** (.0011)	.056*** (.0011)	.076*** (.0022)	.077*** (.0018)
Age ² /100				-.049*** (.00092)	-.053*** (.0012)	-.051*** (.0012)	-.075*** (.0024)	-.078*** (.0020)
Constant	2.94*** (.0036)	2.74*** (.013)	2.65*** (.013)	1.37*** (.021)	1.69*** (.037)	1.71*** (.038)	5.85*** (.055)	5.88*** (.048)
Industry	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Occ dummies							2dgt	4dgt
Observations	155,781	155,781	155,781	155,781	63,925	63,925	63,925	63,925
R-squared	.137	.288	.358	.451	.527	.563	.424	.486

Note: Robust standard errors in parentheses. *** p<.01, ** p<.05, * p<.1. All models are weighted using sampling weights and contain a dummy for the survey year. Industry controls are dummies.

at the 2-digit industry level. The indices for routine and nonroutine work are all standardised by year for the entire population with zero mean, and a standard deviation of one. The coefficients can therefore be interpreted as the relationship between a one-standard deviation change in routine work on the dependent variables, which is the log of wages on most cases. Because sampling weights are included in each regression, the results can also be seen as representative for the Dutch population.

The first estimation shows that routine work is associated with lower wages: a one standard increase in routine-intensity is followed by an approximate decline in hourly wages of 3 percent¹⁵, in absence of any industry and demographic controls. The wage premium for abstract work is 14%. In columns (2) to (4), industry dummies, education controls and other demographic variables are added to the specification consecutively. The signs of the routine and abstract coefficients do not change, but are reduced in size by roughly half.

¹⁵Note that these are log-linear models, where the wages are log transformed. As a result, to interpret the coefficients they should be exponentiated. However, for small coefficients $e^{\hat{\beta}}$ should roughly equal $1 + \hat{\beta}$, and as such the coefficients on tasks $\times 100$ can roughly be interpreted as percentage changes.

Column (5) shows the same model as column (4), but restricted for the sample that has isco-occupations data - to make a comparison possible with column (6). This sample still contains sampling weights in order to make the sample representative for the Dutch population. Column (6) shows the occupation-level routine and abstract scores taken from Autor and Dorn (2013). The coefficients are in the same direction and the model has a similar explained variation as those using individual-level task data. However, they differ from the worker-level task coefficients in their larger size. Also, including occupation-level measures reduces the importance of the education coefficients. Potentially, this is a feature of sorting in skills on the occupation level, that is captured by using occupations as bundles of tasks, rather than individual-level task data.

The results from both using occupation- and individual level data indicate that the presence of abstract tasks seems to be more important than not doing routine tasks: the coefficient size is consistently higher and significantly different from each other. This implies that doing routine work hurts workers less than not doing abstract work, which indicates the importance of nonroutine cognitive tasks in the Dutch labour market.

In addition, in columns (7) and (8) I add 2 and 4 digit occupation dummies respectively. The goal is to see the extent to which occupational averages capture the effect of individual level tasks. If both the worker-level coefficients would turn insignificant, this would make the use of such data superfluous in the presence of occupation-level data. Though this seems to be the case for routine-intensity, abstract-intensity remains positive and significant and shows that, even when accounting for 4 digit occupations, abstract tasks are associated with a 4.3% wage premium. Note that most of this variation is already picked up by adding 2 digit occupations, and the signs for individual-level task efficiencies are not considerably reduced by more disaggregated occupational classifications.

So, adding occupational dummies reduces the size of the task coefficients, most likely due to self-selection into occupations. Whether this selection is based on general ability (absolute advantage) or task-specific abilities (comparative advantage) cannot be concluded from these regressions.

2.5.2 Comparative advantage and technology complementarities

To test whether workers positively or negatively self-select into occupations based on their abilities, I estimate Equation (2.9). Given the cross-sectionality of the data at hand, even these results can only provide suggestive evidence for tasks as ex-

planatory factor of wage inequalities. The goal of the later analysis of switching patterns is to find a way to disclose how people that initially sorted into certain tasks operate on the labour market later on, which gives an idea of whether this is based on pure ability (absolute advantages) or comparative advantage, in that workers move to places where they can maximise their wages based on task-technology complementarities. Workers with high nonroutine task efficiencies should move to technology-rich industries and vice versa. These estimations are necessary to exclude the possibility of high-ability workers sorting into abstract work, which would be more indicative of skill-biased technological change, rather than routine-biased technological change. Therefore, in order to derive conclusions for RBTC, the combination of cross-sectional analyses and switching patterns are necessary for showing that workers behave according to task-related comparative advantages.

Autor and Handel (2013) estimate this equation and find that the routine interaction term is positive and highly significant, whereas the abstract interaction term is negative. In combination with the fact that occupation and worker measures of routine are negatively correlated with wages it suggests that routine tasks are relatively more prevalent in low- and middle-paying jobs. This is potentially consistent with a Roy model, where workers who are productive at a certain set of tasks will self-select into occupations that differentially reward those tasks. I estimate an augmented Mincer regression as in Autor and Handel (2013), with the inclusion of group-level averages $\bar{T}_j$:

$$\ln w_{ij} = \beta_0 + \sum_{T \in (N,R)} (\beta_T T_i + \delta_T \bar{T}_j + \gamma_T T_i \bar{T}_j) + X_{ij} \beta_X + Z_j \beta_Z + \varepsilon_i \quad (2.17)$$

The results from the Dutch data are presented in Table 2.6, with robust and clustered standard errors (at the 2 digit industry level). Columns (1) to (4) include industry averages of tasks, whereas columns (5) to (8) contain occupation-level averages. Again, the sample of the latter 4 models is smaller than in the baseline regressions, because it is restricted for the sample that includes isco-occupation dummies in each specification. Uneven columns present models with only task coefficients, and even columns additionally include interaction terms with group-level and individual-level task measures.

In comparing column (3) and (4) we see that adding industry-level task means does not change the sign of the task coefficients. In the final column, with the full set of controls, we observe that routine work shows signs of comparative advantage, whereas abstract work is not necessarily more valued in industries where abstract

Table 2.6: OLS Regressions of Log Hourly Wages on Worker-level Tasks and Task Averages Grouped on Industry and Occupation

	Task means $\bar{R}_j$ and $\bar{N}_j$ denoting:							
	Industry Averages				Occupation Averages			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
R_i	-.026*** (.001)	-.033*** (.001)	-.013*** (.001)	-.017*** (.001)	-.018*** (.002)	-.030*** (.002)	-.004*** (.002)	-.014*** (.002)
N_i	.106*** (.001)	.111*** (.001)	.0677*** (.001)	.068*** (.001)	.073*** (.002)	.086*** (.002)	.50*** (.002)	.058*** (.002)
$\bar{R}_j$	.169*** (.007)	.182*** (.007)	-.210 (.140)	-.18 (.150)	.12*** (.007)	.14*** (.007)	.060*** (.007)	.068*** (.007)
$\bar{N}_j$	.56*** (.006)	.58*** (.006)	.11 (.120)	.18 (.120)	.52*** (.006)	.53*** (.006)	.32*** (.006)	.33*** (.006)
$R_i \times \bar{R}_j$		.13*** (.004)		.079*** (.004)		.11*** (.004)		.078*** (.004)
$N_i \times \bar{N}_j$		.024*** (.003)		-.011*** (.003)		.039*** (.003)		.024*** (.003)
Controls	No	No	Yes	Yes	No	No	Yes	Yes
Obs.	155,781	155,781	155,781	155,781	63,925	63,925	63,925	63,925
R-squared	.222	.230	.451	.453	.304	.319	.512	.519

Note: Robust standard errors in parentheses. *** p<.01, ** p<.05, * p<.1. All models are weighted using sampling weights and contain a dummy for the survey year.

work is used more intensively. This seems to indicate that abstract work can be implemented across industries, whereas routine work is specifically valued in industries with more routine tasks. Still, performing routine tasks is negatively associated with wages, but the impact is cushioned by working in a routine-intensive industry. The industry-level wage premium partly offsets the negative association between routine tasks and wages.

Next, I introduce occupation-level task means. Column (6) and (8) show that the interaction terms between routine and abstract tasks are positive. This implies that routine (abstract) tasks are valued more in occupations that use routine (abstract) tasks intensively: a sign of sorting based on comparative advantages in tasks. These findings corroborate those of Autor and Handel (2013), who also find positive coefficients on the routine-interaction term (which is consistent with a Roy model of self-selection). However, their results on abstract-intensity show a negative coefficient, whereas we find a positive one - though smaller than for routine. Therefore, using individual-level task data, I find evidence for sorting on comparative advantage for nonroutine abstract tasks on the occupation-level.

With the knowledge of sorting based on comparative advantage, I move to estimate task interactions with ICT to test complementarities between tasks and

technology. These interactions should be indicative of the technological change element in RBTC: does the presence of technology change the estimated task prices? To analyse this, I estimate the following equation:

$$\ln w_{ij} = \beta_0 + \delta_K K_j + \sum_{T \in (N,R)} (\beta_T T_i + \gamma_T T_i K_j) + X_{ij} \beta_X + Z_j \beta_Z + \varepsilon_i + \mu_s \quad (2.18)$$

This model with the inclusion of ICT parameters on the industry level is presented in Table 2.7. Including the full set of demographic controls, the same pattern for the raw task coefficients emerges: not significant for routine and positive for abstract tasks. Column (2), (4) and (6) show two additional findings. First, wages in general are higher in ICT-intensive industries. This should be expected, given the presence of κ_j in the function for marginal product of labour. Second, and perhaps more interestingly, routine-intensity is associated with a stronger wage penalty in ICT-intensive industries. This significance holds, even with the inclusion of 2 digit occupation dummies in column (6). Even though the main coefficient for routine-intensity is insignificant, doing routine work in more ICT-intensive industries is associated with a wage penalty. In other words, routine-intensity is specifically valued less in industries that have replaced routine tasks with ICT. Contrarily, abstract-intensity is not valued more in ICT-intensive industries. This does not provide evidence for the hypothesis that abstract work complements ICT and should therefore see higher returns.

A similar analysis in column (7) and (8) using Autor and Dorn (2013) task data shows a similar pattern as before. However, interestingly, here routine-intensity on the occupation level is not significantly negative with the inclusion of ICT indicators. Furthermore, the interaction term is also insignificant. In other words, within-occupation variation in routine tasks can explain a wage penalty for routine workers in technology rich environments, whereas occupation-level data cannot pick up these effects. Individual-level task data can thus confirm RBTC on the worker-level, whereas this is not picked up by occupation level data. This is the main advantage of using such worker-level data for the analysis of RBTC, especially in a time period where task replacement has been occurring for a significant amount of time.

Table 2.7: OLS Regressions of Log Hourly Wages on Task Intensities and Industry- level Computer Use

	Worker-level Tasks				Occupation Tasks			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Routine	-.031*** (.001)	-.020*** (.001)	-.012*** (.001)	-.007*** (.001)	-.001 (.002)	-.002 (.002)	-.017*** (.002)	-.007*** (.002)
Nonroutine	.14*** (.001)	.13*** (.001)	.081*** (.001)	.077*** (.001)	.056*** (.002)	.056*** (.002)	.17*** (.002)	.16*** (.002)
Industry computer use		.10*** (.001)		.066*** (.001)		.041*** (.002)		.042*** (.002)
Routine × Industry computer use		-.015*** (.001)		-.017*** (.001)		-.015*** (.002)		-.025*** (.002)
Nonroutine × Industry computer use		.009*** (.001)		-.000 (.001)		.003*** (.002)		.001 (.002)
Constant	2.94*** (.004)	2.93*** (.004)	1.33*** (.019)	1.38*** (.018)	1.60*** (.035)	1.60*** (.035)	1.08*** (.024)	1.13*** (.024)
Demographic controls	No	No	Yes	Yes	Yes	Yes	Yes	Yes
Occ. dummies (2dgt)	No	No	No	No	Yes	Yes	Yes	Yes
Observations	155,781	155,781	155,781	155,781	63,925	63,925	79,176	79,176
R-squared	.137	.198	.377	.401	.483	.491	.469	.482

Note: Robust standard errors in parentheses. *** p<.01, ** p<.05, * p<.1. All models are weighted using sampling weights and contain a dummy for the survey year. 2digit occupation (occ) tasks data from O*NET (Goos and Salomons, 2009).

Table 2.8: OLS Regressions of Log Hourly Wage Growth between $t + x$ and t_0 on Task Intensities and Industry-level Averages in Tasks and Computer Use

	t+2			t+5			t+8			Industry average $\bar{I}_j =$		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	$\bar{K}_j$	$\bar{R}_j$	$\bar{N}_j$
Routine	.00036 (.00096)	-.00055 (.00096)	.0025** (.0013)	-.00025 (.0012)	.0027 (.0018)	-.00028 (.0018)	-.00012 (.0012)	-.00046 (.0013)	-.00010 (.0012)			
Nonroutine	.0054*** (.0010)	.0032*** (.0010)	.0093*** (.0013)	.0047*** (.0013)	.014*** (.0019)	.0062*** (.0019)	.0051*** (.0013)	.0053*** (.0013)	.0053*** (.0013)			
Middle educ		-.0017 (.0022)		-.0049* (.0028)		-.0060 (.0043)	-.0047* (.0028)	-.0049* (.0028)	-.0048* (.0028)			
High educ		.018*** (.0025)		.027*** (.0033)		.045*** (.0049)	.027*** (.0033)	.027*** (.0033)	.027*** (.0033)			
Female		.0037*** (.0017)		.0040* (.0022)		.00046 (.0034)	.0042* (.0022)	.0041* (.0022)	.0041* (.0022)			
Age		-.0080*** (.00080)		-.017*** (.0011)		-.020*** (.0017)	-.017*** (.0011)	-.017*** (.0011)	-.017*** (.0011)			
Age ² /100		.0069*** (.00092)		.013*** (.0012)		.015*** (.0021)	.013*** (.0012)	.013*** (.0012)	.013*** (.0012)			
$\bar{I}_j$							.055*** (.021)	.055*** (.021)	.055*** (.021)			
$R_i \times \bar{I}_j$							.0019 (.0013)	.00093 (.0046)	.0052 (.0041)			
$N_i \times \bar{I}_j$							.0032*** (.0012)	-.010** (.0045)	.0083** (.0038)			
Constant	.032*** (.0090)	.24*** (.019)	.026** (.012)	.47*** (.024)	.066*** (.020)	.63*** (.039)	.56*** (.040)	.53*** (.031)	.93*** (.17)			
Observations	102,859	102,859	104,205	104,205	54,067	54,067	104,205	104,205	104,205			
R-squared	.014	.027	.013	.055	.016	.077	.055	.055	.055			

Note: Robust standard errors in parentheses. *** p<.01, ** p<.05, * p<.1. All models are weighted using sampling weights and contain a dummy for the survey year.

2.5.3 Longitudinal results

For the longitudinal analysis, I look at hourly wage growth in a period of 2, 5 and 8 years after filling out the survey. The baseline estimation is the following:

$$\ln w_{ijt+x} - \ln w_{ijt_0} = \beta_0 + \beta_1 R_{ij} + \beta_2 A_{ij} + X_i \beta_3 + \varepsilon_{ijt} \quad (2.19)$$

Column (1) to (6) in Table 2.8 present the results from Equation (2.19), where uneven columns exclude demographic controls and even columns have all controls. All estimations are with robust standard errors. Again, the routine coefficient is non-significant. Abstract work is associated with higher wage growth over a period of 2, 5 and 8 years. This path dependency, where doing more abstract work at time t_0 leads to faster wage growth in the period after, is also something found in other longitudinal analyses of e.g. being in a declining occupation (Edin et al., 2019) or in a routine occupation in the 1980s (Cortes, 2015). The finding implies that wage inequality based on abstract tasks increases over time, leading to stronger disparities over the longer run.

Next, I also estimate a series of these wage growth estimations with the inclusion of industry-level averages, to see whether there is a path dependency created based on whether working in certain industries has an impact on future careers. Here, I look at the average ICT, routine and abstract-intensity of an industry. This leads to the following function for $I \in \{K, R, N\}$:

$$\ln w_{ijt+5} - \ln w_{ijt_0} = \beta_0 + \sum_{T \in (N, R)} (\beta_T T_{ij} + \gamma_I \bar{I}_j T_i) + \delta_I \bar{I}_j + X_i \beta_3 + \varepsilon_{ijt} \quad (2.20)$$

Where $I \in \{K, R, N\}$ are the industry averages. The results are presented in columns (7) to (9) of Table 2.8. Workers who work in a relatively computer- or abstract-intensive industry at the time of the survey see stronger wage growth in 5 years after, where the coefficient for nonroutine-intensity is specifically strong. On the other hand, working in a routine-intensive industry is negatively related to wage growth. The positive relationship is amplified for those workers who perform nonroutine tasks in a computer-intensive or nonroutine industry (see column (7) and (9)). Interestingly, the magnitude of the coefficient for the industry-level task means is larger than that of the individual-level tasks. However, performing nonroutine tasks in a routine tasks dampens the positive relation between nonroutine tasks and wages: even though the nonroutine coefficient remains significant and positive in column (8), performing nonroutine tasks in a relatively routine-intensive industry

negates that positive relationship.

In other words, if workers sort according to comparative advantage in abstract tasks, they are rewarded for this on the medium-long run. If they perform abstract tasks in a routine environment, thus not having sorted according to comparative advantages, they face a wage penalty - that is double in size of the positive coefficient of doing abstract work.

These results thus indicate persisting and growing inequalities between relatively abstract workers and non-abstract workers, which is amplified by the initial routine-intensity of their environment, in this case their industry.

Given the interest in sorting patterns, I am also interested in the switching decisions of workers. More precisely, I aim to estimate the probability of switching to an industry that is either more or less intensive in tasks and ICT in the period of 5 years after the survey. The target industry is the last industry that we observe in the year $t + 5$. I abstract from analysing the amount of switches in this period, but only observe the worker in t_0 and t_5 to see if they have switched. I measure whether this target industry is more or less intensive in tasks or ICT in $t + 5$ than the industry the worker is in at t_0 . This leads to six logit regressions, with six different estimated probabilities: switching either up (u) or down (d), in either ICT K_j , routine R_j or abstract N_j intensity of the target industry.

I do this analysis using three types of dependent variables. In the first six estimations, I use the population-standardised routine and abstract-intensity. These results indicate whether workers who are more routine/abstract than the average Dutch worker are more likely to switch. This is estimated using Equation (2.21). Second, I use the home-industry deviation from the mean as indicator of a 'push' factor in terms of tasks $T_i - \bar{T}_j$ and level of education $E_i - \bar{E}_j$. We should expect that workers who are more routine, abstract or higher educated than their peers are more likely to move to other industries. These estimations measure whether the probability of switching increases when workers are more or less task intensive than their industry-peers. This is presented in Equation (2.22).

Third, I use the difference between the target industry j' and the individual task intensities as indicator for a pull factor. This estimation is presented in Equation (2.23). If the difference between the target industry and the individual tasks is high, this should affect the probability of switching. If the average worker in another industry is more intensive in a certain task, this should work as a pull factor for those currently intensive in that task. The larger this difference is, the more likely a worker

should be to move there in order to maximise their returns to task efficiencies.

$$P[S = 1] = \Lambda(\beta_0 + \sum_{T \in (n,r)} \beta_T T_i + \gamma_E E_i) \quad (2.21)$$

$$P[S = 1] = \Lambda(\beta_0 + \sum_{T \in (n,r)} \beta_T^{push} [T_i - \bar{T}_j] + \gamma_E^{push} [E_i - \bar{E}_j]) \quad (2.22)$$

$$P[S = 1] = \Lambda(\beta_0 + \sum_{T \in (n,r)} \beta_T^{pull} [\bar{T}_{j'} - T_i] + \gamma_E^{pull} [\bar{E}_{j'} - E_i]) \quad (2.23)$$

With S equals a dummy for whether the average worker in target industry is either more or less intensive in the execution of (non)routine tasks. There are four cases for tasks: switch to a higher industry in terms of task-intensity ($S \equiv \tilde{T}_i^h = \bar{T}_{j'} - T_i > 0$, $T \in \{R, N\}$) or to a less task intensive industry ($S \equiv \tilde{T}_i^l = \bar{T}_{j'} - T_i < 0$, $T \in \{R, N\}$). Thus, the switch dummy compares the average task-intensity in the target industry at time $t + 5$ to the task-intensity of the individual worker at time t_0 . For the ICT switch dummy, I compare the average computer use in the target industry to the home industry. Therefore, S can take four values for the measuring the difference in tasks and two for the difference industry-level computer use: $S \in \{K_j^h, K_j^l, R_j^h, R_j^l, N_j^h, N_j^l\}$. The results are presented in Table 2.9 and the six values for S correspond to the six columns in each panel.

Panel A describes the results Equation (2.21) for the standard task scales. The coefficients capture whether workers that are relatively routine as compared to the general population are more likely make certain switches. Columns (1) and (2) show the likelihood of switching to an industry with either more or less average computer use than the home industry, conditional on the tasks in t_0 and the level of education¹⁶. The only significant coefficient for the task scales is for routine tasks, which is associated with lower odds of switching to a less ICT-intensive industry. However, the converse case is not true: workers executing relatively more routine tasks are not necessarily more likely to switch to an industry that uses more computers, as indicated by the nonsignificant coefficient in column (1). In other words, workers do not tend to switch according to complementarity to or substitutability with technology.

Column (3) to (6) in Panel A show results for switching according to comparative advantages in tasks. If such sorting patterns exists, we should see that (non)routine workers switch to industries where the average worker executes more (non)routine tasks than they currently do ($\bar{T}_{j'} > T_{i0}$). Using the population-standardised measures, I find no clear evidence for such patterns. If they switch, workers executing

¹⁶Demographic controls are added in all estimations, but not presented in Table 2.9.

Table 2.9: Logit Regressions of Probability of Switching to Industry that is Higher or Lower in ICT, Routine, or Nonroutine-intensity than Home Industry

A.Task scales, standardised for entire population	ICT		Routine		Nonroutine	
	Higher	Lower	Higher	Lower	Higher	Lower
	(1)	(2)	(3)	(4)	(5)	(6)
R_i	.056 (.038)	-.099*** (.038)	-.013 (.038)	-.030 (.039)	-.23*** (.042)	-.049 (.038)
N_i	-.023 (.038)	-.033 (.038)	.093** (.038)	-.15*** (.038)	.15*** (.041)	.067* (.037)
Middle educ	-.23*** (.085)	.35*** (.087)	.22*** (.085)	-.11 (.087)	.58*** (.097)	.24*** (.085)
High educ	-.32*** (.092)	.42*** (.094)	.059 (.091)	.032 (.094)	1.31*** (.10)	.088 (.092)
Obs	5,008	5,008	5,008	5,008	5,008	5,008

B. Deviation from home industry mean	ICT		Routine		Nonroutine	
	Higher	Lower	Higher	Lower	Higher	Lower
	(1)	(2)	(3)	(4)	(5)	(6)
$R_i - \bar{R}_j$	-.031 (.039)	.019 (.038)	.15*** (.038)	-.17*** (.040)	-.067* (.039)	.050 (.039)
$N_i - \bar{N}_j$	.043 (.038)	-.12*** (.037)	-.049 (.038)	-.027 (.039)	-.021 (.038)	-.13*** (.038)
$Educ_i - Educ_j$	.27*** (.049)	-.22*** (.048)	-.24*** (.048)	.29*** (.050)	.34*** (.050)	-.33*** (.048)
Obs	5,008	5,008	5,008	5,008	5,008	5,008

C. Deviation from target industry mean	ICT		Routine		Nonroutine	
	Higher	Lower	Higher	Lower	Higher	Lower
	(1)	(2)	(3)	(4)	(5)	(6)
$\bar{R}_{j'} - R_i$	-.13*** (.038)	.17*** (.038)	1.00*** (.11)	-.047 (.074)	.090* (.054)	.0061 (.074)
$\bar{N}_{j'} - N_i$	.079** (.038)	-.022 (.039)	-.29*** (.063)	.23*** (.050)	-.017 (.065)	-.21*** (.051)
$Educ_{j'} - Educ_i$	.50*** (.047)	-.50*** (.047)	-.12* (.063)	.31*** (.071)	-.27*** (.075)	-.38*** (.073)
Obs	5,008	5,008	2,911	2,097	2,085	2,097

Note: Robust standard errors in parentheses. ***p<.01, **p<.05, *p<.1. All models include only the set of switchers, are weighted by sampling weights, contain a dummy for the survey year and contain the full set of demographic controls. Samples in Panel C column (3) and (5) only include workers that have a lower task-intensity than the target industry j' , and in column (4) and (6) only a higher task-intensity than the target industry j' in order to test the size of the difference between the two industries in the decision to shift to a higher or lower industry.

relatively more nonroutine tasks are more likely to switch to more routine-intensive industries and to more nonroutine-intensive industries at the same time. A likely explanation is that comparing tasks across the population are less indicative of sorting, than when using push and pull factors as described in equations (2.22) and (2.23).

Panel B therefore shows whether workers who are more (non)routine-intensive than their peers are more likely to switch to a more or less (non)routine-intensive industry. Here, I find that relatively routine workers within an industry are more likely to switch to an even more routine-intensive industry. In other words, they move to an industry where the average worker is more routine-intensive than they are currently. Thus, when workers are relatively more routine than their peers, they are more likely to sort into production processes that require more routine labour. They are thus pushed away from their own industry, into a more routine-intensive industry. The converse is not the case for nonroutine workers: the deviation from the industry mean is not related to significant switching patterns to higher or lower routine-intensive industries. However, they are less likely to switch to a less nonroutine-intensive industry (column 6).

Interestingly, there are clearer patterns for the education variable. Workers that are higher educated than their industry-peers are more likely to switch to industries that use more computers, are less routine-intensive, and more nonroutine-intensive. They are also significantly less likely to switch down in ICT and nonroutine-intensity, and up in routine-intensity. Relatively high skilled workers thus follow technology and nonroutine tasks: they sort to industries that need more workers that can work with technology and can execute creative, problem-solving and learning tasks.

Panel C shows the results on the pull factors: if the difference between the target industry and workers' individual tasks is larger, are they more likely to switch to that industry? The samples in Panel C column (3) and (5) only include workers that have a lower task-intensity than the target industry j' , and in column (4) and (6) only a higher task-intensity than the target industry j' in order to test the size of the difference between the two industries in the decision to shift to a higher or lower industry.

If the distance between the target industry and the person's routine tasks is larger, workers are more likely to switch to a more routine-intensive industry (column 3). The smaller the difference in nonroutine tasks between the target industry's average and the individual, the lower the probability this worker will move to a more routine-intensive industry or less nonroutine-intensive industry. Also, a larger difference is significantly related to a higher likelihood of switching to a less routine-intensive industry. Furthermore, Panel C columns (1) and (2) show that more higher ICT-

intensive industries attract more nonroutine workers, especially if these workers are far less nonroutine than the nonroutine-intensity of the target industry. The converse holds for the deviation between target industry's routine-intensity as compared to the individual.

In summary, the results from Panel A to C show the following pattern. When compared to the workers in their current industry of employment, the relatively routine workers are more likely to switch to more routine-intensive industries. This finding is in line with the results in Table 2.6 which shows that routine tasks have higher returns in more routine-intensive industries. Other sorting patterns on tasks (Panel A en Panel C) do not provide clear evidence for either sorting according to comparative advantage or technology complementarities. However, the education variables are significant in almost all estimations. Relatively high educated workers within an industry tend to behave according to complementarities between skills and technologies: they move to industries that employ more nonroutine labour and use more computers.

For the last longitudinal estimation, I analyse how wage growth of workers differs conditional on having switched or stayed in the same industry. To analyse this, I include a dummy S in Eq. 2.19 that can take the value 1 for staying in the same industry in the period of 5 years after the survey, 2 if switched up, and 3 if switched down to a target industry in terms of ICT, routine or abstract-intensity. Note that this could still imply that workers have switched employers or occupations within the same industry. The results of these estimations are presented in Table 2.10.

From all estimations, I retrieve the same signs and sizes for the routine and abstract coefficients as in Table 2.8. Columns (1) and (2) present the wage growth analysis with the inclusion of the switch dummy in terms of ICT. Here we learn the following. First, switchers earn less than stayers. Workers who stay within the same production process are probably more likely to pick up important tacit knowledge in their field, which is not easily transferable to other industries. Second, switchers who have switched up in ICT-intensity ($2.S_k$) face a lower penalty than those switching down in ICT-intensity ($3.S_k$). It is thus better to switch to a more ICT-intensive industry, if one needs to switch. Third, the interaction terms between the switch dummies and tasks are mostly insignificant, apart from routine-intensity for those who have switched to a less ICT-intensive industry. This coefficient is positive. If routine workers switch, they can best go to an industry where they have to 'compete' with fewer computers in doing routine tasks.

In column (3) and (4) we find similar results. Again, switching is worse than staying. However, here we see that switching to a more routine-intensive industry

Table 2.10: OLS Regressions of Log Hourly Wage Growth between $t + 5$ and t_0 on Task Intensities, Conditional on Switching

	ICT-intensity		Routine-intensity		Nonroutine-intensity	
	No controls (1)	With controls (2)	No controls (3)	With controls (4)	No controls (5)	With controls (6)
Routine	.0021* (.0012)	-.00078 (.0012)	.0022* (.0012)	-.00071 (.0012)	.0020* (.0012)	-.00095 (.0012)
Nonroutine	.0087*** (.0012)	.0037*** (.0012)	.0086*** (.0012)	.0036*** (.0012)	.0084*** (.0012)	.0034*** (.0012)
<i>Ref cat.: no switch</i>						
Switch to (dummy):						
More intensive	-.079*** (.011)	-.094*** (.011)	-.15*** (.0097)	-.16*** (.0095)	-.051*** (.014)	-.074*** (.014)
Less intensive	-.14*** (.010)	-.15*** (.010)	-.060*** (.012)	-.077*** (.012)	-.15*** (.0097)	-.16*** (.0095)
More intensive × Routine	-.0057 (.013)	-.0034 (.013)	.016 (.012)	.019 (.012)	.0095 (.019)	.0096 (.019)
Less intensive × Routine	.028** (.014)	.031** (.014)	.0093 (.015)	.011 (.014)	.017 (.012)	.020* (.012)
More intensive × Nonroutine	-.0064 (.012)	-.0021 (.012)	.0014 (.012)	.0078 (.012)	-.031 (.019)	-.027 (.018)
Less intensive × Nonroutine	.0089 (.012)	.014 (.012)	.0065 (.013)	.0084 (.012)	.0056 (.012)	.011 (.012)
Constant	.031** (.012)	.49*** (.025)	.030** (.012)	.49*** (.024)	.028** (.012)	.51*** (.024)
Observations	101,068	101,068	101,068	101,068	100,265	100,265
R-squared	.022	.065	.023	.066	.023	.068

Note: Robust standard errors in parentheses. *** $p < .01$, ** $p < .05$, * $p < .1$. All models are weighted using sampling weights, contain a dummy for the survey year and 2 digit industries and contain the full set of demographic controls. Reference category are the non-switchers, "more intensive" means a dummy for a switch to a more intensive industry in terms of ICT (columns (1) and (2)), routine-intensity ((3) and (4)) and nonroutine-intensity ((5) and (6)). Interaction terms reflect whether more (non)routinized task intensive workers that switch to a more I intensive industry have faster/slower wage growth than nonswitchers.

Table 2.11: Regressions on Various Indicators for Job Quality

	Job security			Job sustainability			Other	
	Tenure in t_0 (1)	Tenure in $t + 1$ (2)	Unemp in $t + 1$ (3)	Jobs in year t_0 (4)	Log Hours (5)	Full time (6)	Remote work (7)	Job satisf (8)
Routine	.058*** (.015)	.085*** (.016)	.068*** (.014)	.041*** (.0079)	.012*** (.0016)	.0041 (.0085)	-.16*** (.014)	-.43*** (.0081)
Nonroutine	-.18*** (.014)	-.21*** (.015)	-.12*** (.015)	-.063*** (.0081)	.058*** (.0018)	.26*** (.0089)	.39*** (.017)	.094*** (.0084)
Middle educ	-.065** (.031)	-.034 (.035)	.029 (.033)	-.0043 (.019)	.021*** (.0037)	.039** (.020)	.30*** (.039)	-.073*** (.019)
High educ	-.33*** (.039)	-.39*** (.046)	.13*** (.038)	.083*** (.021)	.053*** (.0043)	.34*** (.023)	1.04*** (.039)	-.023 (.022)
Female	.082*** (.26)	.025 (.29)	.085*** (.24)	-.0022 (.13)	-.28*** (.031)	-1.85*** (.16)	-.22*** (.26)	.11*** (.15)
Observations	155,781	153,721	153,721	155,781	155,781	155,780	90,766	155,265

Note: Robust standard errors in parentheses. *** $p < .01$, ** $p < .05$, * $p < .1$. All models are weighted using sampling weights, contain a dummy for the survey year and 2 digit industries and contain the full set of demographic controls. Column (1) to (4) are poisson regressions. The dependent variable in column (1) and (2) are the number of months spent in a non-tenured contract in t_0 and t_1 respectively, for column (4) the number of months spent unemployed in the year after the survey (not in t_0 because the working conditions survey is aimed towards only the working population). Column (4) measures the number of jobs (i.e. employers) in the year of the survey. Column (5) is an OLS regression on the log of monthly hours. Column (6) to (8) are logit regressions. The dependent variable in column (6) is a dummy that takes the value 1 when a worker works full-time (i.e. more than or equal to 36 hours) and 0 if not. The dependent variable in column (7) is a dummy that takes the value 1 if the worker can work remotely, and 0 if not, and column (8) the dummy takes the value 1 if a worker is either satisfied or very satisfied, and 0 if not.

is affecting wage growth more negatively than switching to a less routine-intensive industry. Again, there is a positive interaction for routine workers moving to a more routine-intensive industries: sorting on routine comparative advantage can mitigate the negative effect of switching to a routine-intensive industry. Lastly, column (5) and (6) present the results for wage growth conditional on switching to a more or less abstract-intensive industry. Again, both are negative, but moving to a more abstract-intensive industry is less negative than to a less abstract-intensive environment. Also, moving to a more abstract-intensity bears to smallest negative in the table: this is the least adverse switch to make. And again, the results show that routine workers are better when they move to a less-abstract-intensive industry ($3.S_a$). This is thus a consistent pattern that shows up in all three groups of estimations.

2.5.4 Additional job quality indicators

The conceptual framework accommodates for inequality in hourly wages between workers of different task groups. In this subsection, I highlight that returns to labour are not only monetary in nature, and can also be expressed in terms of non-pecuniary outcomes. Such measures in job quality indicate the willingness of employees to

invest in their workers, e.g. by providing them with a tenured contract, more contracted hours, or possibilities to work from home. This also bears interest because job insecurity can have significant negative physical and mental health implications for workers (De Witte, Vander Elst, and De Cuyper, 2015). Besides, it creates a risk of downward spiraling as insecure workers have intentions to quit, show reduced commitment, reduced life and job satisfaction and lower self-rated performance (Ashford, Lee, and Bobko, 1989; Cuyper and Witte, 2006).

In this section, I show that other job quality measures, though non-pecuniary in nature, also reflect growing inequalities in between routine and abstract workers. Following Weil (2014), these may be important indicators of the growing fissure between lovely and lousy jobs (Goos and Manning, 2007). The findings presented in Table 2.11 indicate that routine and abstract tasks are related to this fissure: routine-intensity is structurally related to worse indicators of job quality and stability, whereas abstract tasks are associated with better quality jobs.

From Table 2.11 we learn the following. First, routine workers are more likely to have a non-permanent contract (columns 1 and 2), both in t_0 as well as in the year after the survey. Furthermore, they have more unemployed months (column 3) and a higher number of jobs in the year after the survey (column 4). The opposite is true for abstract work. Besides that, abstract work is associated with more contracted hours, and they are more likely to work full-time (columns 5 and 6). In light of the necessity of working from home in times of COVID-19, I have also included column (7) to show that abstract work can more easily be done from home, whereas routine work is more associated with working on a specific location, and not with teleworking. Lastly, column (8) shows that job satisfaction is far lower when workers have performing one-standard deviation more routine tasks than the average worker. A positive, but less strong, coefficient is found for relatively abstract work.

2.5.5 Summary of results

The findings presented above lead to the following summary of results, see Table 2.12. All in all, abstract work is not only associated with higher wages and higher wage growth, it is also clearly related to better job quality. The results highlight the fact that the inequalities on the current labour market seems to be characterised by routine-biased technical change.

In summary, I find that a one standard deviation in abstract tasks is significantly associated with the following: i) higher wages and better job quality, also when controlling for up to 4 digit occupation dummies, ii) even higher wages when

Table 2.12: Empirical Results for i) worker-level measures, ii) industry-level averages of tasks and ICT and iii) interaction terms.

Estimation	Eq.	Worker-level				Industry-level				Interaction terms			
<i>Cross-sectional</i>													
Baseline	(2.8)	β_N	+	β_R	n.s.								
Comparative advantage	(2.9)	β_N	+	β_R	-	δ_N	+	δ_R	n.s.	γ_N	n.s.	γ_R	+
ICT interactions	(2.10)	β_N	+	β_R	n.s.	θ_K	+			ϑ_N	n.s.	ϑ_R	-
<i>Longitudinal</i>													
Baseline	(2.12)	β_N	+	β_R	n.s.								
Industry interactions: $\bar{N}_j$	(2.13)	β_N^L	+	β_R^L	n.s.	δ_N^L	n.s.			γ_N^L	+	γ_R^L	n.s.
Industry interactions: $\bar{R}_j$	(2.13)	β_N^L	+	β_R^L	n.s.			δ_R^L	-	γ_N^L	-	γ_R^L	n.s.
Industry interactions: $\bar{K}_j$	(2.13)	β_N^L	+	β_R^L	n.s.	δ_K^L	n.s.			γ_N^L	n.s.	γ_R^L	n.s.

Note: Check marks indicate coefficients were as predicted, for a significance of up to 10%. n.s. is non-significant.

workers are relatively abstract within a relatively abstract occupation, iii) faster wage growth over a period of at least 8 years, iv) even higher wage growth if the worker started out in an abstract-intensive industry, but lower when started in a routine-intensive industry, v) higher probability of switching to a more abstract-intensive industry, vi) lower probability of switching to a less abstract or ICT-intensive or more routine-intensive industry.

On the contrary, a one standard deviation in routine tasks is associated with i) lower wages (but only in absence of demographic or occupational dummies) and significantly worse job quality ii) higher wages when working in a relatively routine-intensive industry or occupation, iii) lower wages in ICT-intensive industries, iv) lower probability of switching to a more abstract-intensive industry, v) mitigating effect of wage decline in response to switching to a more routine, less ICT or less abstract-intensive industry (thus when sorting according to comparative advantage)

2.6 Discussion

The findings presented above are consistent with a theory of nonroutine biased technical change in the 21st century in the Dutch labour market. Workers performing abstract tasks have higher wages, which exceed the wage penalties for routine work.

In order to put the results from the Netherlands into an international perspective, I have rerun the analysis on cross-sectional wage inequality using the European Working Conditions Survey (Eurofound, 2019; Parent-Thirion, Biletta, Cabrita, Llave Vargas, Vermeylen, Wilczynska, and Wilkens, 2017). The results and data construction are presented in the appendix. I find similar results for the Dutch sample in the EWCS, which suggests that the patterns I document extend beyond

the Netherlands. Again, nonroutine tasks are associated with higher wages, and routine tasks are insignificant, controlling for 2 digit occupational data. For many countries, the routine estimate is insignificant, in contrast to a positive and significant coefficient on abstract-intensity in almost all countries. Very similar results can be obtained using occupation-level task data: there is a premium for nonroutine tasks, which at times is also accompanied by a negative wage return to routine tasks. From this exercise, I conclude that the Dutch labour market is not unique, but operates in line with inequalities that exist across the European continent.

The results from the analyses that interact the ICT indicator with worker-level tasks could be extended using different measures of ICT. The ICT analysis in its current form focuses on the average computer use in an industry, and thus indicates the direct interaction of workers with technology. The measure excludes the possibility of technology adopted in the workplace that is not used by workers directly, e.g. software that runs in the background. In order to solidify the results on the interaction between tasks and technology, adding different technological parameters would be an interesting extension to this line of work. For instance, the EUKLEMS database provides data on technology stocks and flows for computing equipment, communications equipment and computer software and databases separately. Furthermore, Statistics Netherlands also provides possible ICT indicators on the industry-level data from firm-level surveys. This consists of data on e.g. access and use of internet, percentage of workers working with a computer, different types of software used, innovations in new processes or products, use of big data analysis or investments in ICT personnel. Another possibility is to interact tasks with an entire different type of technology, for instance by the inclusion of robot data from the International Federations of Robots.

Furthermore, the task data used in this paper stems from a working conditions survey. This implies it may also be used for different countries that adopt similar surveys. The European Working Conditions Survey is the largest and most prominent of these surveys. However, working conditions surveys have also been conducted in the USA (Maestas, Mullen, Powell, Von Wachter, and Wenger, 2017), South Korea (Park and Lee, 2009), and several Latin American countries such as Colombia, Argentina, Chile, Uruguay and a group of countries in Central America (Merino-Salazar, Artazcoz, Campos-Serna, Gimeno, and Benavides, 2015). Especially when combined to (local) register data, these provide interesting uses for the micro-level analysis of RBTC for multiple countries.

2.7 Conclusion

Using job-level information on tasks, this paper contributes to the literature by showing how tasks vary within occupations, and how this variation shapes careers, both in terms of monetary rewards, other measures of job quality, wage growth and switching patterns. Specifically, this paper is the first to use a Dutch working conditions survey matched with longitudinal administrative information to study these issues. The Dutch data are particularly opportune for studying the labor market consequences of routine and abstract work because they cover a substantially large sample, allowing for variation at the most detailed occupation level, for the most recent period. I show that the Dutch labour market can still be characterised by RBTC.

The results highlight underlying task-dynamics of wage polarisation. Routine-biased technical change can explain the displacement of routine-intensive occupations in the 80's, 90's and early '00s. Using worker-level data, I show that such dynamics are still apparent in the labour market of the 2000's and 2010's: nonroutine work is associated with significant wage premiums. Even though routine work is not significantly related to wages and thus wage inequality across the entire sample, routine tasks are associated with wage differences and switching patterns that follow a theory of sorting on comparative advantage. Results underlining this finding are that i) routine work is valued more in industries that use routine labour more intensively, ii) workers who perform relatively more routine tasks tend to switch to more routine-intensive industries, and iii) wage declines are less strong for routine workers that switch to either more routine-intensive, or less abstract and ICT-intensive industries.

A dichotomy on the labour market can thus still be found. In the Dutch 21st century labour market, this is characterised by a strong association between abstract work and jobs with decent wages, permanent contracts, job stability and job satisfaction. Routine tasks are associated with lower job quality, and seem to be forced to move to areas where routine work is still prevalent. In other words, they do not tend to move to more abstract or ICT-intensive industries, but rather move to the places where routine work is still prevalent. And these industries also provide better wages than industries with lower routine-intensity.

The results are not only informative for the understanding of the underlying micro-dynamics of job polarisation, but they are also instructive for policy makers aiming to harness the benefits of automation and mitigate its losses. This necessitates a better understanding of the nature of jobs, and how it relates to individual outcomes.

Moving beyond occupations, the paper shows that person-level job characteristics can be explanatory for future wage growth. Providing more training opportunities for workers that enables them to perform abstract tasks could be worthwhile to ensure that the future of work is meaningful and dignifying.

2.A Survey questions

Dutch Working Conditions Survey (NEA)

Abstract

- Can you decide how to execute your work? (v05h_a)
1) Yes, often; 2) Yes, sometimes; 3) No
- In your work, are you required to think of solutions to do things? (v05h_d)
1) Yes, often; 2) Yes, sometimes; 3) No
- Does your job require you to learn new things? (v05m_b)
1) Never; 2) Sometimes; 3) Often; 4) Always

Routine

- Are you doing work where you have to make repetitive movements? (v05d_c)
1) Yes, often; 2) Yes, sometimes; 3) No
- Can you decide the sequence of your tasks? (v05h_b)
1) Yes, often; 2) Yes, sometimes; 3) No
- Can you decide the speed of your tasks? (v05h_c)
1) Yes, often; 2) Yes, sometimes; 3) No

European Working Conditions survey (EWCS)

Abstract

- Are you able to choose or change your methods of work? (Q54b)
1) Yes, 2) No, 8) Don't know, 9) Refusal
- Generally, does your main paid job involve solving unforeseen problems on your own? (Q53c)
1) Yes, 2) No, 8) Don't know, 9) Refusal
- Generally, does your main paid job involve learning new things? (Q53f)
1) Yes, 2) No, 8) Don't know, 9) Refusal

Routine

- Please tell me, does your main paid job involve repetitive hand or arm movements? (Q30e)
1) All of the time, 2) Almost all of the time, 3) Around 3/4 of the time, 4) Around half

of the time, 5) Around 1/4 of the time, 6) Almost never, 7) Never, 8) Don't know, 9) Refusal

- Are you able to choose or change your order of tasks? (Q54a)
1) Yes, 2) No, 8) Don't know, 9) Refusal
- Are you able to choose or change your speed or rate of work? (Q54c)
1) Yes, 2) No, 8) Don't know, 9) Refusal

Table 2.A.1: Comparing occupational means in task-intensity from NEA and EWCS to O*NET

ISCO	O*NET			NEA			EWCS		
	R	A	sd	R	A	sd	R	A	sd
12 Corporate managers	-.98	1.69	.66	-.52	.55	.60	-.67	.61	.53
13 General managers	-.98	1.69	.74	-.45	.31	.78	-.63	.58	.59
21 Math, engineering, physical science prof.	-.68	1.40	.78	-.34	.54	.62	-.34	.47	.76
22 Life science and health professionals	-.02	2.36	.95	-.06	.37	.69	.23	.20	1.00
24 Other professionals	-1.41	1.19	.82	-.33	.42	.75	-.49	.49	.69
31 Physical & engineering science assoc prof.	.32	.80	.93	-.16	.34	.78	-.09	.14	.91
32 Life science and health associate prof.	.34	.27	1.02	.06	.20	.83	.61	-.33	1.12
34 Other associate professionals	-1.16	.65	.84	-.29	.32	.80	.08	.07	.96
41 Office clerks	-1.09	-.49	.94	-.09	.01	.95	-.21	-.28	.86
42 Customer services clerks	-.65	-.43	1.05	.44	-.27	1.03	.34	-.46	1.15
51 Personal and protective services workers	-.02	-.45	1.02	.12	-.09	.99	.45	-.45	1.10
52 Models, salespersons and demonstrators	-.76	-.60	1.04	.15	-.29	1.07	.26	-.47	1.08
71 Extraction and building trades workers	1.06	-.31	.97	.05	.05	.96	.26	-.11	.91
72 Metal, machinery and other trades workers	1.23	.34	1.04	.15	.04	.97	.24	-.05	.95
73 Trades workers	.89	-1.36	.09	.09	.01	1.01	.53	-.38	1.04
74 Other craft and related trades workers	.77	-1.41	.38	.38	-.34	1.12	.07	.18	.88
81 Stationary-plant and related operators	1.39	-.56	1.10	.41	-.30	1.14	.82	-.75	1.05
82 Machine operators and assemblers	1.37	-.53	1.08	.71	-.52	1.13	.57	-.83	1.10
83 Drivers and mobile-plant operators	1.39	-.66	1.06	.56	-.51	1.02	.90	-.80	1.05
91 Sales & services elementary occupations	.02	-1.44	1.00	.05	-.60	1.04	.26	-.92	.96
93 Labourers in manufacturing etc.	.63	-1.07	1.05	.66	-.77	1.18	.80	-1.11	1.06
Correlation Goos and Salomons (2009)				.72	.84		.75		.78

Note: Data from Goos and Salomons (2009) and author's calculations using NEA and European Working Conditions Survey as extra comparison (see Appendix 2.B). European occupational averages standardised on European level, not within-country standardisation. R stands for routine, A for nonroutine abstract.

Table 2.A.2: Dividing the population by occupation and typology

	Job Typology			Total
	Abstract	Routine	Mix	
ISCO Level 1 <i>Simple and routine tasks, Low or unskilled</i>	1.9%	.3%	2.8%	5.0%
ISCO Level 2 <i>Little to medium complex tasks, Middle skilled</i>	24.3%	2.8%	1.9%	38.0%
ISCO Level 3 <i>Complex tasks, Middle or high skilled</i>	15.4%	1.1%	3.5%	2.0%
ISCO Level 4 <i>Very complex tasks, High skilled</i>	31.5%	1.1%	4.4%	37.0%
Total	73.1%	5.4%	21.6%	100%

2.B The Netherlands in the European context: results from EWCS

In order to put the Dutch results in a broader perspective, I rerun similar analyses using the European version of the Working Conditions Survey (EWCS). The goal of this exercise is threefold. First, it creates an extra validation of the task-based measures used in this paper. Second, it allows me to corroborate the cross-sectional results on wages from the Netherlands. Third, I can compare countries in their routine-intensity and the effect it has on labour market outcomes, and thus provide an extra analysis of underlying dynamics of RBTC.

I use data from the European Working Conditions Survey. The EWCS collects information every 5 years for a number of countries in Europe, including the Netherlands. I use data from 2005 to 2015 to match the same time period as the Dutch data. The EWCS has the same questions regarding routine and abstract tasks. To validate the task measures used in the Dutch data, I also construct the same measures for the European data. For a comparison of the questions, refer to Appendix 2.A. A summary of the task indices in comparison to the Dutch data and O*NET data are presented in table 2.A.1 in the Appendix.

Descriptive statistics of the European data are provided in the Appendix Table 2.B.1. The EWCS provides data for 33 countries for the years 2010 and 2015, for samples of approximately 1,000 to 3,000 workers per country. The total sample contains 53,186 workers, of which half is female. 61% of the European workers can be categorised as Abstract, versus 14% as Routine and the remaining 25% as a Mix.

For the employment characteristics, the EWCS only provides self-reported monthly wage. Therefore, there is no analysis possible for future employment.

Figure 2.B.1 shows all countries and their respective routine intensities, in a similar fashion as Figure 2.2. Countries are ranked by their abstract-intensity (bottom to top). The Netherlands is located in the bottom of the figure, together with the Nordic countries, plus Malta and Estonia. Compared to the European mean, these countries are both more abstract-intensive and less routine-intensive. Bulgaria, Greece, Romania and Croatia top the list of most routine-intensive countries. To infer to what extent these countries are routine-intensive in absolute terms, the right pane is informative. Less than half of Bulgarian workers have abstract-intensive jobs, against over 80% of Norwegian, Danish and Maltese workers. The mix category is roughly equal over the countries, such that an increase in routine-intensive workers is typically accompanied by a lower share of abstract workers.

I also estimate equation (2.8) using the European data. The results are provided in Table 2.B.2, and for ease of comparison, I have depicted coefficients for routine and abstract-intensity of each country in Figure 2.B.2. All estimations include 2 digit occupation dummies, which takes up the largest part of self-selection. Column 1 shows the result for the total sample: the abstract task scale has a positive coefficient, whereas the routine task scale is, though statistically significant, not so in economic terms. Similar to the results from NEA, for the Netherlands I find that there is a positive effect of abstract intensities, and a non-significant effect of routine-intensity (see Column 19).

Figure 2.B.3 shows the results where the individual-level measure in (2.8) is substituted by the occupation-level measure by Goos and Salomons (2009). Again, the same pattern emerges: abstract-intensity has a larger effect than routine-intensity, which is hardly significant in any country. Therefore, using occupation-coding results in the same conclusion on the relative importance of abstract tasks. Nevertheless, differences emerge as the ranking of countries in Figure 2.B.2 deviates from that in Figure 2.B.3. More research should be devoted to understanding these differences, and whether this would fundamentally change our view of the impact of technology on labour markets.

Table 2.B.1: Descriptive statistics for European sample

		Total			Types		
		Total	Men	Women	Routine	Abstract	Mix
General							
Observations		53,186	26,286	26,900	7,630	32,393	13,163
Age			(49%)	(51%)	(14%)	(61%)	(25%)
Education		43	43	44	43	44	43
	Low	.20	.23	.17	.31	.17	.22
	Middle	.47	.49	.46	.56	.43	.53
	High	.32	.28	.37	.13	.40	.25
Female							
Current Job							
Log monthly wage		6.78	6.90	6.66	6.37	6.93	6.64
		(.94)	(.95)	(.92)	(.85)	(.95)	(.89)
Occupation							
	(1) Managers	.06	.09	.05	.01	.10	.03
	(2) Professionals	.18	.14	.22	.04	.23	.14
	(3) Technicians and assoc. prof	.13	.12	.14	.06	.15	.12
	(4) Clerical support workers	.09	.06	.12	.08	.09	.10
	(5) Service and sales workers	.20	.14	.26	.25	.18	.22
	(6) Skilled agri, forestry, fishery	.03	.04	.02	.03	.04	.03
	(7) Craft and trades workers	.12	.20	.04	.15	.11	.13
	(8) Plant and machine operators	.08	.12	.03	.18	.04	.10
	(9) Elementary occupations	.10	.09	.12	.20	.07	.13
Tasks							
	Routine-intensity (standardised)	.00	.01	-.01	1.38	-.62	.71
	Abstract-intensity (standardised)	.01	.04	-.01	-1.58	.51	-.28

Note. Data from EWCS.

Figure 2.B.1: Comparison of routine-intensity between countries, in relative terms using PCA and in absolute terms using the typology

Source: Author's calculations using EWCS data.

Table 2.B.2: OLS Regressions of Log Monthly Wage on Task Scales, Demographic Controls and Occupation Dummies by Country

	(1) TOT	(2) BEL	(3) BGR	(4) CZE	(5) DNK	(6) DEU	(7) EST	(8) GRC	(9) ESP	(10) FRA	(11) IRL	(12) ITA	(13) CYP	(14) LVA	(15) LTU	(16) LUX	(17) HUN
Routine	.0051** (.0025)	-.016** (.0065)	-.0021 (.015)	-.051*** (.019)	-.0014 (.011)	-.022 (.015)	.027 (.017)	.00055 (.015)	.050*** (.013)	.010 (.010)	.0100 (.017)	-.012 (.014)	-.0094 (.015)	.021 (.017)	.023 (.019)	.010 (.016)	-.0039 (.018)
Abstract	.055*** (.0026)	.044*** (.0077)	.078*** (.017)	.024 (.016)	.052*** (.016)	.065*** (.016)	.089*** (.018)	-.053*** (.015)	.083*** (.013)	.046*** (.011)	.067*** (.018)	.053*** (.014)	.066*** (.015)	.013 (.018)	.048** (.021)	.045*** (.016)	.013 (.019)
Middle educ	.10*** (.0069)	.067*** (.019)	.24*** (.044)	.16*** (.039)	.0056 (.047)	.18*** (.074)	.10*** (.038)	.072* (.042)	.080*** (.027)	.048 (.032)	.16*** (.052)	.072** (.032)	.077* (.040)	.051 (.048)	.15** (.060)	.088** (.036)	.070 (.051)
High educ	.28*** (.0087)	.24*** (.021)	.49*** (.062)	.47*** (.078)	.059 (.058)	.38*** (.092)	.31*** (.043)	.14*** (.044)	.33*** (.037)	.22*** (.040)	.25*** (.060)	.21*** (.045)	.22*** (.050)	.31*** (.061)	.35*** (.070)	.36*** (.049)	.40*** (.075)
Female	-.25*** (.0054)	-.20*** (.016)	-.20*** (.034)	-.31*** (.046)	-.18*** (.035)	-.39*** (.037)	-.31*** (.038)	-.19*** (.032)	-.29*** (.027)	-.21*** (.024)	-.37*** (.036)	-.26*** (.027)	-.24*** (.030)	-.24*** (.039)	-.27*** (.046)	-.27*** (.032)	-.085** (.037)
Obs	52,916	3,713	1,387	1,277	1,639	1,573	1,153	1,146	2,712	3,386	1,395	1,424	1,180	1,385	1,382	1,018	1,115
R-squared	.712	.304	.394	.273	.215	.359	.443	.394	.279	.289	.352	.356	.349	.341	.404	.450	.243

	(18) MLT	(19) NLD	(20) AUT	(21) POL	(22) PRT	(23) ROU	(24) SVN	(25) SVK	(26) FIN	(27) SWE	(28) GBR	(29) HRV	(30) MKD	(31) TUR	(32) NOR	(33) ALB	(34) MNE
Routine	.00056 (.012)	-.0021 (.015)	.0072 (.018)	.024 (.016)	.017 (.015)	.046*** (.014)	.0076 (.0099)	.00016 (.014)	-.010 (.012)	-.0073 (.012)	-.0015 (.017)	.014 (.015)	.035*** (.016)	.024** (.011)	.0055 (.011)	-.049*** (.017)	.022 (.015)
Abstract	.027** (.013)	.083*** (.019)	.078*** (.018)	.050*** (.017)	.042*** (.014)	.046*** (.015)	.061*** (.010)	.040*** (.013)	.048*** (.012)	.037*** (.013)	.076*** (.019)	.039** (.017)	.037** (.016)	.094*** (.012)	.036*** (.012)	.014 (.021)	.035** (.016)
Middle educ	.14*** (.035)	-.014 (.047)	.087 (.061)	.24*** (.067)	.12*** (.034)	.16*** (.050)	.16*** (.042)	.22*** (.093)	.012 (.046)	.067 (.045)	.030 (.038)	-.043 (.032)	.20*** (.043)	.20*** (.024)	.052 (.052)	.079* (.043)	.15** (.072)
High educ	.26*** (.042)	.12** (.052)	.25*** (.076)	.50*** (.079)	.42*** (.063)	.45*** (.063)	.45*** (.048)	.39*** (.10)	.076 (.052)	.18*** (.056)	.17*** (.042)	.13*** (.047)	.44*** (.062)	.38*** (.033)	.15** (.058)	.29*** (.054)	.42*** (.082)
Female	-.32*** (.032)	-.43*** (.038)	-.37*** (.033)	-.19*** (.037)	-.16*** (.029)	-.12*** (.033)	-.17*** (.023)	-.28*** (.029)	-.13*** (.030)	-.14*** (.028)	-.35*** (.034)	-.17*** (.034)	-.15*** (.034)	-.23*** (.024)	-.23*** (.024)	-.22*** (.040)	-.20*** (.032)
Obs	1,180	1,528	1,400	1,309	1,273	1,262	2,186	1,274	1,647	1,517	1,898	1,361	1,331	2,767	1,644	1,286	1,168
R-squared	.311	.352	.337	.280	.406	.426	.367	.352	.300	.233	.416	.340	.295	.332	.338	.188	.355

Note: Robust standard errors in parentheses. *** p<.01, ** p<.05, * p<.1 All regressions are with survey year controls (2010 and 2015), as well as industry, age and 2 digit occupation controls. Task scales are standardised by country. A constant is included in each model, but not included here.

Figure 2.B.2: Coefficients of Abstract and Routine task scales from OLS Regressions of Monthly Wages for European Countries.

Source: Author's calculations using EWCS data. All regressions include demographic controls and industry and 2 digit occupation dummies. Task scales are standardised by country. Gray dots indicate confidence intervals

Figure 2.B.3: Coefficients of Abstract and Routine occupation-level task scales from OLS Regressions of Monthly Wages for European Countries.

Source: Author's calculations using EWCS data. All regressions include demographic controls and industry and 2 digit occupation dummies. Gray markers indicate confidence interval.

2.C Data construction explanation

2.C.1 Matching NEA to POLISBUS

- **Wages:** For the wage data I use the variable SLNINGLD, which are the total gross wages received in a given month. It is the wage indicator over which taxes are calculated. I create the hourly wage variable by dividing this variable by the SBASISUREN, which are the total monthly contracted hours.
- **Wage growth:** I calculate wage growth by calculating the average hourly wage in a given year $t + \tau$, and from that subtract the average hourly wage in the year t_0 , as in the following equation where m denotes a month.

$$\Delta \ln y = \ln \left(\frac{1}{12} \sum_{m=12} \left(\frac{slningld_m}{sbasisuren_m} \right) \right)^{t+\tau} - \ln \left(\frac{1}{12} \sum_{m=12} \left(\frac{slningld_m}{sbasisuren_m} \right) \right)^{t_0}$$

- **Switching:** For the switching, I create a dummy taking a value of 1 if the 2 digit industry (sbi) is not equal to the sbi in 5 years after the survey.
- **Unemployment:** The matched POLISBUS data is an unbalanced panel if unemployed months is not included. I use the STATA command `xtfill`, which takes a value of 1 if a person is not present in any POLISBUS file in that month, and thus is unemployed. I sum the number of 1's over a x year period to obtain the total number of months unemployed. I replace the value by 0 if a person is above 60 at the time of the NEA, and never returns in the POLISBUS: I will then assume that the person retired. This is not a watertight procedure (as people may have become self-employed or for other reasons leave the sample), and the estimates on unemployment should therefore be interpreted with some caution.
- **Temporary contracts:** I use the variable SARBEIDSRELATIE, which is a dummy labelled 1 if a person has a 'vast' (tenured) contract and a 0 if a person has a 'flex' (flexible) contract (either uitzendkracht or oproepkracht).

3

Learning the Right Skill

3.1 Introduction

Technology, globalisation and structural change have worsened labour market outcomes of middle educated workers, and especially for the relatively recent entrants to the labour market. Even though this has been related to changing demand in skills and tasks, we know very little about the skills that middle educated students are currently learning in school, and the returns to those skills. This chapter addresses two questions: (i) how focused are curricula on certain skills? and (ii) what are the returns to those skills? Based on recent literature, we focus on social, technical and basic cognitive skills. We develop a new measure of the relative weight of these three skills in curricula, by extracting text from the full set of training curricula in the Dutch middle education system. By linking the curriculum data to register data on wages for the graduates, we estimate skill returns on a fine-grained, curriculum-content level. To the best of our knowledge this is the first that uses the granulated skills level, using curriculum data, to highlight the importance of these skills for graduates entering the labour market.

The central notion behind this chapter originates from the literature describing how technology has created poorer labour market trajectories of middle skilled workers. New technologies and offshoring have caused a decline in the labour demand for routine tasks, which are historically often executed by middle-skilled workers (Autor et al., 2003; Autor and Dorn, 2013; Goos et al., 2014; Goos and Manning, 2007). This is accompanied by increased demand for complex, analytical and social tasks, which are more complementary to technology (Autor et al., 2006; Spitz-Oener, 2006). At the same time, there is a general upskilling of the workforce, where more workers than ever are highly educated (Goldin and Katz, 2010). This upskilling has also led to higher-skilled workers taking up less skilled occupations than before, increasing the average skill level within occupations (Beaudry and Green, 2003; Spitz-Oener, 2006). These features of the contemporary labour market make it increasingly difficult for middle-skilled workers to have fulfilling careers: routine occupations are disappearing, the remaining occupations are becoming more complex, and the set of occupations that do match their skill level are increasingly

This chapter is based on the GLO Discussion paper no. 979 “Learning the Right Skill: The Returns to Social, Technical and Basic Skills for Middle-Educated Graduates” (Cnossen, Piracha, and Tchuente, 2021) and is currently a Revise & Resubmit at an international journal. The authors are indebted to Robert Inklaar, Milena Nikolova, Steven Brakman, Noemi Peter, Juliette de Wit, Olena Nizalova, Klaus Zimmermann, Christian Siegel, and participants from the GLO Virtual Young Scholar program, ESPE (2021), and SOM Conference (2021) for helpful comments and advice. We would specifically like to thank Max Boiten, Elwin Buisman and Welmoed Cnossen for support related to the text analysis, and the Stichting Beroepsonderwijs Bedrijfsleven for access to the curriculum data. Any errors are our own. Personal contribution to this chapter: idea creation, data collection, data analysis, writing.

filled by higher skilled workers.

Even though automation increases high skill job creation in the service sector, it reduces middle skill job creation in the manufacturing (goods) sector (Dauth et al., 2021). Students can adjust by either attending college to increase their skill level, or to take jobs requiring more abstract and less routine intensive tasks. However, this is not an option for all students. In the Netherlands, this is reflected by a decreasing probability of employment for middle-educated students, and especially if they are trained for routine professions (Bisschop, Zwetsloot, Ter Weel, and Van Kesteren, 2020; Ter Weel et al., 2021).¹ Yet, the middle education sector in the Netherlands still accounts for over 40% of the student population after high school.²

Despite the decline in middle-skilled employment, Autor (2015) expects that a ‘significant stratum of middle-skill jobs combining specific vocational skills with foundational middle-skill levels of literacy, numeracy, adaptability, problem solving, and common sense will persist in the coming decades’ (p. 27). He also points out that this prediction strongly depends on the education system being able to teach the current generation of middle-educated workers the “right” skills. This notion is the key motivation for this chapter. Rather than a focus on one-dimensional skill, it makes a case for shifting towards an analysis of multi-dimensional skills. In this chapter, we allow students from the same level of education to differ in the type of skills they have been taught in schools, in order to estimate skill-based wage inequalities in a novel way.

The extant literature shows a number of specific skills that are increasing in importance. The most prominent are the so-called people skills; there is growing evidence of relative employment and wage growth for occupations that require social skills (Borghans, Ter Weel, and Weinberg, 2014; Deming, 2017). These are jobs that require high levels of coordination, persuasion, negotiation, social perceptiveness, influencing, and decision-making (Felstead, Gallie, Green, and Zhou, 2007; Borghans et al., 2014; Deming and Kahn, 2017; Deming, 2021). Relevant to this chapter, Deming (2017) highlights the importance of the combination of skills. More specifically, jobs where social skills are combined with high levels of cognitive skills have fared well, which is shown both using occupational task data (Deming, 2017) as well as using data from job postings (Deming and Kahn, 2017). In contrast, the opposite happened to high-math, low-social skills jobs (including many Science, Technology, Engineering

¹See also Reinhold and Thomsen (2017) who show that German students graduating from middle education have seen decreasing starting wages and slower wage growth than cohorts before the turn of the century.

²Source: CBS Statline, accessible via <https://opendata.cbs.nl/statline/#/CBS/nl/dataset/03753/table?fromstatweb>.

and Math – or STEM – occupations).

However, most of the current literature on the technology-induced changes in demand for different types of skill focus on the higher-earning end of the labour market. For instance, the findings of Deming and Kahn (2017) are based on a sample restricted to professional occupations that employ predominantly college-educated workers. This is because vacancy data (in this case Burning Glass) has the most representative coverage for this group of occupations. Hansen, Ramdas, Sadun, and Fuller (2021) use only a sample of executive occupations to show the growing importance of social skills. Following Deming (2017), the growing importance of social skills in high-paying occupations might be explained by the fact that there is specific complementarity between cognitive and social skills. Therefore, solely improving middle-educated students' social skills cannot be seen as a panacea for potential substitution if it is not accompanied by other skills, or strong cognitive abilities.

Besides social skills, there is also evidence for a need of skills related to working with specific technologies. STEM graduates earn more than other graduates in the first years after graduation (Deming and Noray, 2018), which can be explained by the fact that STEM graduates have technology-specific skills that complement technologies in a unique way. It is still unclear whether this is caused by the actual skills students are learning, as there is little insight into the curricula of STEM students. Nevertheless, it is very plausible that the technological orientation of these programs increase complementarity between their taught skills and a labour market characterised by fast technological change. And again, the main sample of these analyses are college-educated workers, whose patterns need not be the same for middle-educated graduates.

There is a large stream of literature on the changing demand for task-specific skills, with a number of methods of measuring this change in demand. Such measurements might involve: i. increases in the employment or wages of occupations with certain task-intensities (e.g. Autor and Dorn, 2013), ii. the growth or decline of certain activities within occupations (e.g. Spitz-Oener, 2006), iii. the change in return to specific tasks within occupations (e.g. Chapter 2 of this dissertation) or iv. the change in skills described in vacancies (Deming and Kahn, 2017). Yet, there is only a small number of papers that describe how skill acquisition at school affects labour market outcomes, and how different sectors vary in the relative labour demand for skills in recent graduates.

We use data from the Foundation for Cooperation on Vocational Education, Training and Labour Market (SBB) to measure the skills that are described in the

training curricula of the Dutch middle education system. We merge the curriculum data to non-public microdata from Netherlands Statistics to obtain labour market outcomes of its graduates based on the skills taught. More precisely, we first measure skills from the curriculum text by extracting verb noun combinations from skill descriptions. We retrieve the underlying structure of the skills data using both exploratory factor analysis and a labeling of verb noun combinations to three classes of generic skills: social, technical and basic skills. Second, we link the skills data to the labour market outcomes of the 322,205 students that graduated in the period 2010 to 2018 in the programs included in our curriculum analysis.³ We estimate Mincer (1974) equations of wages in the first years after graduation on our skill measures. We specifically differentiate between the three main fields of education (STEM, economics and healthcare), three middle-educated levels (ISCED level 2, 3 and 4), and apprenticeship-based tracks versus class-based training track.

We present three main findings. First, we find that graduating from relatively social-skill intensive degrees is negatively associated with wages in the first year after graduation while technical skills are associated with positive returns. Both relations persist until at least 10 years after graduation. Second, we show that demand for technical, social and basic skills differ strongly across fields, levels and tracks of education. For instance, students that graduate in the health-related field of education have higher returns to technical skills, as compared to STEM and economics graduates. Third, we show that wage returns to skill are conditional on the sector of employment: social skills are more strongly negatively associated with wages in the high skill service sector than in the low skill service sector. Our results imply that, within the same field of education, degrees focusing relatively more on social skills have lower wage returns. This does not necessarily mean that the demand for social skills is lower for middle-educated students. As our analysis concerns a study of curriculum texts, we show that a relative focus in the curriculum on social skills, rather than on other types of skills, does not positively affect wages. This could have potential implications for the construction of curricula.

Even though education data tends to suffer from endogeneity due to students selecting in degrees based on ability and schools changing curriculum based on local labour market preferences, the nature of our data can partly circumvent these issues. We argue that under three, not highly restrictive assumptions, our estimates can be viewed as consistent. First, students select into degrees, and not in verb noun combinations. It is likely that they choose a field (e.g. STEM) or sub-field (Craft,

³This is the sample of graduates for which we have data for all (control) variables. The actual number of graduates in the middle education system is slightly higher.

Laboratory and Health Technology) based on their own interests and knowledge about abilities. However, upon entering a specific program, students do not know the exact contents of that curriculum, and do not know the differences between degrees within these (sub)fields. This assumption can be substantiated by empirical studies on how Dutch middle-educated students choose a degree: the contents of a curriculum are rarely the reason (Fouarge, Künn, and Punt, 2017). Since this is the source of variation in our data, we deem it safe to assume that conditional on choice of field, the skills acquired by students are exogenous. Second, we are aided by the fact that curricula are constructed on a national level: schools have little freedom in changing curricula based on local preferences. Combined with the fact that over 80% of students live with their parents, the set of local skills available are exogenous to the student, and due to the nationally oriented curricula, also to the local labour market.

Unlike the existing literature on vocational and middle-educated schooling systems (e.g. Malamud and Pop-Eleches, 2010; Golsteyn and Stenberg, 2017; Hanushek, Schwerdt, Woessmann, and Zhang, 2017; Eggenberger, Rinawi, and Backes-Gellner, 2018), our focus is not on the specific versus general skills but rather the type of skills that students learn. More precisely, we distinguish between social, technical and basic cognitive skills (such as reading and mathematics) and estimate the returns to these skills.

In terms of methods, this chapter fits into a growing literature in economics that uses text data as main source of information (Gentzkow, Kelly, and Taddy, 2019). For instance, Eggenberger et al. (2018) analyse the contents of training curricula of Swiss middle-educated graduates to measure skill specificity. Deming and Kahn (2017) use keywords in job postings to measure skill demand, while Hansen et al. (2021) map texts from occupational description to skill clusters, in cognitive, interpersonal and operational dimensions. Webb (2019) compares the overlap in verb noun combinations between occupational descriptions and patent data, to estimate the potential of replacement of workers by technology. In this chapter, we perform a content analysis of training curricula, similar to Eggenberger et al. (2018). We use verb noun combinations, similar to Webb (2019), and label these combinations in social, technical and basic skill categories, similar to Deming and Kahn (2017) and Hansen et al. (2021). Furthermore, we extend the analysis by estimating a factor model that retrieves the underlying structure of the extracted verb nouns.

Figure 3.1: Visual representation of the Dutch education system.

3.2 Institutional background: the Dutch mbo

The Dutch middle education (mbo) system is similar in nature to other education systems with distinct vocational education pathway (e.g., German, Swiss). It is oriented to provide vocational education training, where each major (or degree, or program) is linked to a profession. It is similar in nature to junior college education. Figure 3.1 shows the flow-chart of the Dutch education system. The students have a number of choices after they finish their primary education at the age of 12. Our focus is on the blue boxes, which are part of the Senior Secondary Vocational Education (mbo), in the middle of Figure 3.1.⁴ After completing high school, students select into one of the four different levels, each more complex than the other, each having a broader and deeper bundle of skills than the other. Level 1 (the entry-program) is focused on acquiring basic learning and executive skills. This level does not lead to a starting qualification. Consequently, most students use Level 1 schooling as a stepping stone for further vocational training rather than as an entry to the labour market. Level 2 consists of basic vocational training and lasts between 2 and 3 years. Level 3 programs last for 3 to 4 years, and focus on learning to work independently. Lastly, students can enroll in level 4 programs, that also last 3 to 4 years. This level covers middle management training, and prepare students on having leadership

⁴Mbo is an abbreviation of the Dutch name, middelbaar beroepsonderwijs, for the middle vocational education system.

positions in sub-teams. Besides a higher difficulty in cognitive skills, these programs also focus more on responsibility, whereas level 2 and 3 are relatively more focused on foundational skills.

The entry requirements of the levels in the mbo are directly linked to the levels in the preparatory secondary education (pre-mbo levels 1 to 4; see Figure 3.1) or on previously obtained levels within the mbo. Once a degree is obtained, students are free to continue learning, or start earning. If they decide to continue learning, they have two options: i) stay in the same field, but move one level up (skill deepening), or ii) switch fields (skill broadening). The student receiving training for employee fast service can thus choose to either start working, or train to become an assistant supervisor or switch fields. More than two third of the Level 1 students (69%) and around 60% of the Level 2 students continue with their studies while around 40% of the levels 3 and 4 students stay in education (Centraal Bureau voor de Statistiek, 2016). The percentage of students continuing studying to higher education has been declining over the past decade, from more than 40% to 35% for the most recent cohort (Centraal Bureau voor de Statistiek, 2018). In our analysis, we focus on the highest obtained degree in the mbo. Students who continue studying, for instance by going to an applied university, are excluded from the sample.

When selecting a degree, students can choose between two different pathways: either they opt for class-based training (BOL) or apprenticeship-based training (BBL). BOL has a focus on (theoretical) schooling, where roughly 20% of training time is spent as an apprentice. BBL is more oriented towards apprenticeships: its students are required to work at least 24 hours a week for a local firm (roughly 60% of training time). Both orientations lead to the same certification. In our analyses, we distinguish between these two groups of students, to see whether certain types of skills are more valued when taught in apprenticeships or in a school setting. For instance, it is likely that technical skills have higher returns when they are included in an apprenticeship training, as the extra practice with using technologies at work should increase the skill-level of these students.

Each level- and field-specific program has its own unique training curriculum: a qualification file. The Dutch Organisation for Vocational Training and Labour Market (S-BB) cooperates with the mbo schools and representatives from various industries to construct training curricula. They have a legal task (through the Dutch Act on Adult and Vocational Education) in developing and maintaining the entire qualification structure (SBB, 2021).⁵ The total set of qualification files is the main data

⁵The Dutch vocational education system is not characterised by having central examination, but rather a centrally described curriculum for each degree. Schools may decide for themselves how they examine

source for this project.

A last element of the Dutch secondary vocational education system, is that its schools are specifically oriented at their local labour market, which is why they are often referred to as regional education centers. Each school works in close cooperation with local companies that provide apprenticeships. Schools may choose which training curricula to offer, based on their own analysis of what their local labour market needs and what local students want to study. Students have a large degree of freedom in selecting their preferred training curriculum.⁶ However, characteristic of the student population in the middle education system is that they tend not to move for schooling: 80% of students live with their parents during their entire degree (Fouarge et al., 2017). As a result, most students select a degree from the set of training curricula that their local school offers. Furthermore, there is empirical evidence that students hardly take into account the labour market prospects of each degree, but rather choose degrees based on personal interests and abilities, plus the opinions of their friends and families (Fouarge et al., 2017).

3.3 Empirical strategy

The goal of this chapter is twofold. We propose a new measure of granulated skills, and we estimate returns to social, technical and basic cognitive skills in middle educated graduates. In this section, we explain how we develop our measures of skill and how we use these measures for estimating returns to skills.

3.3.1 Measuring skills in curricula

We obtain curriculum data from the S-BB, which collects all qualification files for all degrees in the Dutch middle education system. Each curriculum consists of a list of core tasks, associated with the profession for which the student is trained. Each task is linked to a number of skill descriptions, which are deemed necessary for the execution of each task. We construct our skill measure based on these skill descriptions, see Figure 3.2.

Each curriculum consists of a list of skills, each a sentence long. The average number of sentences per degree is 18.74, with a standard deviation of 13.29. From these sentences, we extract verb noun combinations using a language programming

the students in the prescribed learning goals and competencies from the curriculum. To ensure that the quality of examination does not create significant differences between the skill acquisition of graduates across schools, the Dutch Inspectorate of the Education closely monitors the teaching syllabuses and examination (Ministerie van Onderwijs, 2013).

⁶Note that the actual contents of each degree are decided at a national level.

Figure 3.2: Composition of training curricula.

module for Dutch, called Frog (Van den Bosch et al., 2007). We use two features of this module: first, Frog replaces all words in each sentence by its dictionary form (lemmatizing) and second, the Part of Speech (POS) tag function categorizes each word as either a noun, adjective, verb or other grammatical form. We keep all the verbs and nouns in each sentence, and use these two Parts of Speech to create a verb noun combination.

In terms of cleaning and preparing text data, we make a few selections. First, in order to make sure that the extracted verb noun combination relates to the main skill in each sentence, we only use main sentences, and delete all prepositional sentences. For example, any part of a sentence following prepositions like "such as", "among which", "for example", or "with respect to" are not part of the analysis. Second, we do not take into account sentences that start with "has knowledge of", which only retrieves 'has knowledge', and also captures knowledge-abilities rather than skills. Lastly, the word "is able to" ("can") is deleted, because it is the first verb in each sentence, e.g., "the student is able to ...". This cleaning process ensures that the verb noun combinations that are extracted from the skill sentences capture the central element of skill in each sentence.

Table 3.1 shows an example for all types of sentences used for the verb noun extraction. We let a verb noun combination exist conditional on the fact that there are at most two verbs and three nouns in the sentence. In this way 15% of sentences are deleted, because there is no clear match between the verbs and the nouns in that sentence. In all other cases, all verbs and nouns in the sentence are matched. The full list of extracted verb noun combinations contains 4450 unique sets. To refrain from overly profession-specific verb noun combinations, we restrict our sample to the verbs that exist in at least 5 study programs - which results in a sample of 482 verb noun combinations.⁷

One of the benefits of the Dutch language in this study is the fact that adjectives

⁷We perform robustness analyses using a sample of 152 verb noun combinations that appear in at least 10 curricula, to test sensitivity of our results for sampling verb noun combinations that are more commonly used.

Table 3.1: Example of verb noun extraction, type of sentence, and frequencies of the type of sentences in entire data set

Translated raw text data	Extracted verb noun combinations	Translation	Type	Freq.
Can apply safety requirements	(toepassen, veiligheidsvoorschrift)	(apply, safety requirement)	1 verb, 1 noun	40%
Can use (digital) relevant registration systems and ICT applications	(gebruiken, registratiesysteem) (gebruiken, ICT-toepassing)	(use, registration system) (use, ICT-application)	1 verb, 2 nouns	19%
Can interpret work order cards, drawings or models	(interpreteren, werkorderkaart) (interpreteren, tekening) (interpreteren, model)	(interpret, work order card) (interpret, drawing) (interpret, model)	1 verb, 3 nouns	9%
Can read and understand product information	(lezen, productinformatie) (begrijpen, productinformatie)	(read, productinformation) (understand, productinformation)	2 nouns, 1 verb	8%
Can cooperate and consult with colleagues and supervisors when upholstering furniture	(samenwerken, collega) (samenwerken, leidinggevende) (overleggen, collega) (overleggen, leidinggevende)	(cooperate, colleague) (consult, colleague) (cooperate, supervisor) (consult, supervisor)	2 verbs, 2 nouns	5%
Can read and interpret specifications, drawings and contract documents	(lezen, bestek) (interpreteren, bestek) (lezen, tekening) (interpreteren, tekening) (lezen, contractdocument) (interpreteren, contractdocument)	(read, specification) (interpret, specification) (read, drawing) (interpret, drawing) (read, contract document) (interpret, contract document)	2 verbs, 3 nouns	4%
			Total	85%

are often compounded with a noun, which creates more context for our analysis. An example is presented in the first row of Table 3.1: the extracted noun is "veiligheidsvoorschrift", which translates to "safety requirement". "Voorschrift" (requirement) is the noun in this case, but because of the added information from the compound, it becomes clear that this skill deals with a specific requirement related to safety ("veiligheid"). Similar cases are e.g. ICT-application ("ICT-toepassing"), contract document ("contractdocument") or work order card ("werkorderkaart"). These compounds add contextual information that allows us to improve the understanding of the skill described in each verb noun combination.

There are two options to impose structure on this type of skill data: manual labeling of verb noun combinations and exploratory factor analysis to uncover the underlying structure of the verb noun data. For the former, the researcher chooses to pre-impose a structure, by taking a list of skills currently used in the literature (Deming and Kahn, 2017; Deming and Noray, 2018; Hansen et al., 2021). The latter takes the entire list of skills, and through a factor analysis on the entire sample, matches certain verb noun combinations to each other in a single factor. We adopt both measures here. Though true to the actual structure of the data set, the disadvantage of exploratory factor analysis is that the combination of elements within factors do not always immediately lead to an intuitive overarching theme,

which manual labeling does.⁸

We use the labels for each verb noun combination to construct weighted skill measures for the three main categories. For each skill category, we create a dummy variable that we assign the value 1 if a verb noun combination can be labelled to that skill category (e.g. "cooperate colleague" is a social skill), and a 0 if not. Table 3.3 presents a list of examples, and how these verb noun combinations are matched to categories. We remove duplicates, such that a verb noun combination that might be mentioned multiple times in a curriculum, still counts as one verb noun combination. Using these inputs, we construct the relative frequency of these skills within the curriculum of each degree. The skill-frequency of skill s in degree j is then calculated as:⁹

$$\text{SkillFreq}_{sj} = \frac{\text{Number of verb noun combinations in degree } j \text{ assigned to skill } s}{\text{All verb noun combinations matched to degree } j} \quad (3.1)$$

For the factor analysis, we use a sample of 152 verb noun combinations. We have also performed the factor analysis using the list of 482 verbs, like above. However, for this sample we end up with 75 factors with an eigenvalue higher than 1, of which many factors have empty or too small (<0.4) loadings on verb noun combinations. Especially many of the later factors in this model are simply picking up degrees, rather than underlying structures of skills. Selecting a model with fewer factors did not solve this issue, as the first few factors contain the bulk of verb noun combinations and all other factors load on one or two verb noun combinations. This proved to be a poorer factor model than using the list of 152 verb noun that exist in at least 10 programs, rather than 5. A scree plot of a factor analysis for this sample is presented in Figure 3.A.1a and Figure 3.A.1b.

We choose a model of 36 factors, which is at the cut-off point in the scree plot where the eigenvalue is 1 (see Figure 5a). We rotate using oblique rotation, which assumes that underlying factors may correlate. Given that after a standard varimax

⁸However, manual labeling of verb noun combinations is more prone to researcher bias. Therefore, we also introduced four independent researchers to the data. We presented them the list verb noun combinations, plus a list of O*NET skill descriptions (National Center for O*NET Development, 2021). See <https://www.onetonline.org/find/descriptor/browse/Skills/> for these skills and their descriptions. We asked them to label each of the verb noun combinations to a specific skill on the list, and none if they felt no skill matched the verb noun combination perfectly. As the O*NET skills fall into broader categories skills, we checked whether all connected O*NET skills from these researchers matched to the larger category (i.e. social, technical or basic skills).

⁹Where the maximum value for the denominator is 482: the sample size that contains the most common verb noun combinations extracted and labelled. In the robustness analyses, we also estimate regressions based on a larger sample of 152 verb noun combinations, that appear in at least 10 programs, to show our results are not sensitive to the inclusion of less common verb noun combinations.

rotation, there existed many correlations between the factors higher than 30%, the choice for oblique rotation is more appropriate.

To simplify terminology, for the labeling of each of these factors we again try to closely resemble O*NET occupational skill descriptions (National Center for O*NET Development, 2021). The main factors, their descriptions, and underlying verb noun combinations are presented in Table 3.A.1. There are four main categories of skills that emerge from our data, which do not necessarily capture all skills that O*NET describes. Most skills are technical in nature, such as “management of material resources” (factor 2), which captures elements such as maintaining materials, products and tools. Other skills are social, such as “coordination” (factor 1) or “persuasion” (factor 28). In the category of basic skills, factors emerge such as “reading and following instructions” (factor 8).

The skill-frequency measures from (3.1) and the 36 factors from the factor model are the main independent variables in our estimations on skill returns later on. The first measures the relative importance of certain types of skill in comparison to the total set of described verb noun combinations. The second tries to uncover underlying skills, by retrieving factors. The former gives more intuitive results in our wage estimations, whereas the latter is mostly useful for providing insight in the types of skills that are taught in different fields and levels of education, and how e.g. manufacturing degrees differ from health or economics degrees.

3.3.2 Estimation strategy: returns to skills

In order to estimate the returns to skill, we run Mincer (1974) type regressions, where we regress hourly wages in the year after graduation on our skill measures. The goal of these analyses is to obtain a pattern of revealed skill demand for certain generic skills or competencies, and whether these skill returns differ across fields, levels or type of education. Furthermore, we aim to understand how these patterns might be explained by sector-sorting after graduation, by measuring whether students who have learned certain skills are more likely to be employed in certain industries.

An important point of discussion in the interpretation of the results is the potential risk of selection on skills. First, abler students might self-select into majors with high returns, which would be reflected by high and positive point estimates for certain skills. Besides ability, preferences and interests that influence major-choice might also correlate with labour market outcomes: diligent students are often diligent workers (Arcidiacono, 2004; Altonji et al., 2012). Positive point coefficients in Mincer equations could then wrongfully be interpreted as the return to that skill, rather than

it being a return to the general ability or preference of the student selecting into this skill. This difficulty also explains why there is little hard evidence on the effect of a curriculum on labour market outcomes (Altonji et al., 2012). In part, this is caused by a limited availability of curriculum data on a large scale. However, this gap can mostly be attributed to the fact that student selection into curricula is not a random process.

Even though returns to skill are never completely free from selection bias, we argue that the fine-grained nature of our text data circumvents part of the self-selection problem. Our reasoning is as follows: students will select into a higher hierarchical level (majors) than our observed data (verb noun combinations). It is also one level more disaggregated than e.g. courses, which would be more likely to be part of the student choice (Altonji et al., 2012). In our data, neither students nor schools have any influence over the curriculum requirements, as these are decided on a national level. Majors are thus essentially fixed bundles of skills.

The structure of the data and the institutional setting therefore allows us to partly circumvent selection issues. First, it is likely that students have a preference for learning either social or technical skills. Therefore, we should expect sorting into *fields* of education based on ability and comparative advantage in skills. However, given that students are unaware of the exact contents of the specific degree that they have chosen, our level of observation (verb noun combinations) is exogenous to the students. They cannot control the curriculum, as it is formed nationally, and they also cannot know the relative importance of social, technical and basic cognitive skills beforehand, as they do not access this type of information in their schooling choice. Combined with the fact that 80% of students do not move out of their parents' hometown, also makes sorting on ability less likely: students take the degrees presented to them at their local school as the only given options (Fouarge et al., 2017). Second, the nationally-oriented curriculum also makes it difficult for individual schools to change the curriculum based on their own preferences or local labour market demands. Furthermore, we can also be assured that the skills described in the curriculum will be tested at school, as schools are inspected on compliance with the national curriculum by the Ministry of Education. We can thus assume that each student will acquire the skills mentioned in the curriculum, and we can assume that students do not know in advance which skills they will learn precisely. As these verb noun combinations are the level of observation in our data, this should improve the reliability of our coefficients in light of potential self-selection issues.

Nevertheless, it might be the case that some skills are overrepresented in difficult

programs, i.e. programs that cost more effort for students with low abilities (Deming and Noray, 2018). The correlation between the presence of a certain skill and the average ability within a degree might then be high, which would imply that the skill returns still reflect ability, and thus self-selection. This is something we cannot directly solve with the data at hand. However, this would only be problematic for our results if either all social, technical or basic verb noun combinations that are part of the skill-frequency equation (3.1) would be more difficult than all verb noun combinations in another category. In other words, it would be a concern if, for example, all social verb noun combinations are more difficult than all basic combinations. It is a reasonable assumption to make that this is not the case in our data, and therefore the skill frequency measure would not be a proxy for ability.

Still, the fact that there should be non-random selection of students into majors that might influence the returns to some verb nouns requires us to introduce a series of controls. First, we highlight the importance of including gender and immigration background in our estimations, as they are significant predictors of study field choice. Women and students with an immigration background are underrepresented in STEM related degrees, both in the Netherlands (de Koning, Gelderblom, Den Hartog, and Berretty, 2010) as well as in other countries (MacPhee, Farro, and Canetto, 2013). Reasons for this can be related to academic self-efficacy, where women and students with an immigration background tend to be less confident in finishing a STEM degree - irrespective of their actual academic performances (MacPhee et al., 2013). It has also been associated to group-related preferences or poorer labour market information (de Koning et al., 2010). Whichever the reason of sorting may be, we indeed observe large differences in the student population across fields: students with a migration background are overrepresented in economics degrees, women are overrepresented in health degrees, and both are underrepresented in STEM.

Furthermore, school-level sorting may be an issue in choice of majors. Arcidiacono (2004) shows that high-quality schools make lucrative majors more attractive, and, since high-quality schools attract high-ability students, they contribute to the ability sorting across majors - within school. However, there is little empirical evidence to believe that Dutch middle education students choose majors based on the quality of schools, since they mostly sort into their local school (Fouarge et al., 2017). However, it might still be the case that, within schools, students with certain abilities or preferences choose higher returning degrees. We would want to account for the fact that students within schools are therefore not independently and identically distributed (i.i.d). This is why we cluster all our estimations at the school level, which contains 73 clusters: one for each regional education center. Furthermore, we include

school fixed effects, by adding school dummies in all our regressions.

Besides selection, another empirical question relates to the underlying cause of the relationship between skills and wages. If it is the case that technological differences have an effect on differences in skill demand, we should see different skill returns across sectors of employment. So far, the literature shows that this might be the case, especially when observing changes in the occupational composition in sectors, which reflects a change in the demand for tasks (and thus skills) executed by workers. For instance, the goods sector has seen a strong reduction in the amount of routine workers, whereas the high skill service sector employs relatively far more abstract workers than before but sees no change in the routine-intensity of the average worker (Bárány and Siegel, 2020).

The complementarity between skills and the tasks workers perform in certain sectors might then relate to wage differences between sectors, conditional on skills. To test whether skill demand is equal across the labour market, or whether different production structures require different skill inputs, we perform a few sector-specific analyses. Some skills might have higher returns in more in high-skilled service industries, manufacturing, or low skilled services. Therefore, besides the standard wage equations, we also estimate the relationship between skills and earnings conditional on being employed in a certain sector.

To summarize, estimating wage equations with skill variables poses a risk of endogeneity. In this case, our beta's might reflect sorting on ability, rather than returns to skill. We argue that our estimates can be viewed as consistent, under the following assumptions. First, students select into a field or domain of their preference, but the actual skill acquired in their training is exogenous to that decision. Students are not aware of the relative skill intensity in their programs upon entering a degree. Therefore, our measures of social, technical and basic skills, which are based on verb noun extractions, are exogenous to the sorting decision of the student. Second, we assume that the ability to acquire social, technical and basic cognitive skills is the same. In other words, technical skills are not necessarily less costly for more able students. Third, as each curriculum is constructed at a national level, we have exogenous variation in skill supply. Schools cannot control the curriculum, and thus cannot match the skill supply to skill demand in local labour markets through changes in the curriculum. This is an extra level of exogeneity that improves the consistency of our estimates.

3.4 Data

This chapter relies on two main data sources: curriculum text data from the qualification files of the Dutch VET system and a linked employer-employee data set on earnings and employment. We obtain labour market data of graduates from non-public microdata from Netherlands Statistics. In this section, we present a brief description of the our constructed skills data, and we explain how we link the curriculum-level data to register data on wages and employment.

3.4.1 Curriculum contents: what are students learning?

We can describe the contents of the Dutch curricula in vocational education, based on the skill frequencies and the factor analysis discussed above. Figure 3.3 presents the skill frequencies for the three categories over domains, which is one hierarchy lower than fields. As could be expected, programs in the STEM domains focus more on technical skills, whereas those programs in the field of economics and health are more oriented towards social skills. However, even within the fields we see variation: the STEM domain of Craft, Lab and Health technology is oriented more towards social skills than other STEM programs.

We are also interested in how curricula differ in the relative prevalence of the retrieved factors. These 36 factors are presented in Appendix Figure 3.A.2 by field of education, and a more detailed description over domain for a few selected factors in Figure 3.A.3. In terms of skill factors by sector, the first main distinctive feature of STEM programs is a stronger focus on technical skills. This is of course in line with the nature of these degrees, as they are mostly related to manufacturing and thus are taught more technical skills relative to the other sectors. The social skill that is disproportionally part of STEM curricula is ‘coordination’, which is a factor concerning verb noun combinations related to functioning in a firm, distinguishing one’s specific tasks and responsibilities and being able to deliver reports in meetings.

Economics and healthcare graduates have a more pronounced focus on social skills. However, economics degree have high factor scores for ‘conversation techniques’, ‘sales’ and ‘discussing calculations’. This last factor combines mathematics with cooperation skills. Furthermore, economics students are more likely to learn ‘foreign languages’ and ‘complex problem solving’. Healthcare students are relatively focused on social skills related to managing teams. ‘Management of Personnel Resources’ captures elements such as feedback, reflection and observation skills, combined with instructing others, acting as a contact and recognizing problems. ‘Management in an Office Setting’ is more oriented towards meeting and presenting

Figure 3.3: Standardised skill frequencies, by domain

Source: Authors' calculations using non-public microdata from Netherlands statistics. Note: Skill frequencies calculated following Equation (3.1). The domain 'Safety and Sports' has two bars, as it contains both economics-programs as well as health-programs.

skills, and applying ICT. Apparently, both are dominantly present in such programs.

Figure 3.A.3 outlines how factors differ across domains. The figure presents 4 factors: equipment maintenance (technical), conversation techniques (social), complex problem solving and interpreting manuals (both basic). From Figure 3.A.2 we could already observe that the technical skill of equipment maintenance is specifically dominant in STEM degrees, but here we see that this is mainly because the skill emerges in two domains: construction and process industry. On the other hand, the social skill of conversation techniques and both basic skills have more variation across domains, as well as fields of education. Some STEM domains use conversation techniques intensively, such as mobility and vehicles, but the skill is more often found in the economics degrees, such as administration, catering, and ICT.

3.4.2 Wage data

To estimate returns to skills we link the curriculum data to non-public microdata from Netherlands Statistics on graduation, earnings and employment of students enrolled in these programs. We are able to link each graduate to their respective degree, through unique degree-codes.¹⁰

We construct a linked employer-employee data set using various data sets from the Dutch microdata: wage data in the years after graduation, demographic characteristics (gender, age, migration background), firm-level data to determine the industry of employment, and data on enrollment and graduation from middle and higher education. We use students that graduated between 2010 and 2018 in one of the 333 training curricula.¹¹ Some of these students have been through multiple programs, of which we select the most recent degree. Furthermore, we check whether students did not enroll in a new study program in either middle education or higher education. Those that continued studying and are either still in education in 2018 or have graduated in higher education afterwards are removed from the sample.

The wage data consists of monthly information on hours worked, gross and net wages, type of contract¹², plus industry and employer. We construct an hourly wage

¹⁰Here we use data on graduates (mbo gediplomeerden). This contains data on crebo codes (the Central Register on Professional Education), which we use to link curriculum text information to graduates.

¹¹The full set of curricula contains 500 programs, but we remove programs from the agricultural sector (as these contain a small number students per program) and programs from level 1, as they do not lead to a basic qualification. Furthermore, we only keep full time degrees.

¹²This can be either tenured or temporary, where a tenured contract applies to workers with a contract for an indefinite period, plus interns, directors/major shareholders ('directeur-groootaandeelhouder or dga in Dutch), and people employed under the Sheltered Employment Act (wsw in Dutch). Temporary contracts apply to temporary employees, sub-contracted or on-call employees (uitzendkracht and oproepkracht in Dutch, respectively)

indicator, which is the average hourly wage across multiple jobs in case a worker has more than one job.¹³ We take the yearly average of this hourly wage in the year after graduation as our main dependent variable. Besides wages, the microdata also provides information on the industry of employment.¹⁴ We only use the first digit industry code, which we use for our analysis on the differential demand for skills across sectors. We divide industries into three sectors following Bárány and Siegel (2020): low skilled services (LSS), goods and high skilled services (HSS). Table 3.A.2 describes the classification of industries into sectors, as well as the relative employment and standardised skill frequencies across industries.

Table 3.2 provides descriptive statistics of our main sample, and for subsamples by gender and general field of education. In total, we have 322,205 students in our sample that have information on all control variables, of which most graduated in economics (40%), followed by health (36%) and STEM (24%). Hourly wages are highest for STEM graduates, and lowest for economics graduates. Furthermore, STEM students tend to be somewhat older upon graduation. In terms of gender distribution, STEM programs are highly skewed towards male students, with only 19% female students, whereas 81% of Health graduates are female. Economics degrees have the highest share of students of non-Dutch descent. Lastly, females - and thus health degrees - are over-represented in level 4 programs, whereas STEM has the highest share of level 2 students. Given that STEM is a male-dominated field, this is also reflected in the relative share of men in level 2 programs.

3.5 Returns to curriculum skills

As a first descriptive step in the analysis of the data we perform individual regressions on the entire set of verb noun combinations on wage in the first year following graduation. Wage is the one-year average of log hourly wage, for the first full year after graduation. In Table 3.3 we report the ten verb noun combinations associated with the highest (positive and) significant coefficient for different subsamples.¹⁵ The ranking of verbs is based on regressions with the inclusion of one verb noun combination at a time, with a full set of controls and standard errors clustered at the

¹³This variable is computed using data on number of paid hours and total wage in monetary value, which excludes parts of the wage bill that are transferred in tangible ways, such as lease cars or company lunch.

¹⁴The industry of employment is obtained by matching the employer-identifier we obtain from the POLISBUS data to firm-level register data (ABR - Algemeen Bedrijven Register), which contains industry data. The industry classification used by Statistics Netherlands is the SBI, which has the same first digits as the NACE classification used by the European Union.

¹⁵This type of analysis is inspired by a similar analysis of measuring the importance of certain tasks in metropolitan areas over time by Michaels, Rauch, and Redding (2019).

Table 3.2: Descriptive Statistics

	(1) Total	By gender		By field of study			By sector of emp.		
		(2) Male	(3) Female	(4) STEM	(5) Econ	(6) Health	(7) LSS	(8) Goods	(9) HSS
<i>Wage</i>									
Log hourly, $t + 1$	2.23	2.24	2.22	2.29	2.16	2.27	2.17	2.28	2.29
sd	(.42)	(.40)	(.44)	(.37)	(.41)	(.46)	(.36)	(.42)	(.47)
<i>Demographics</i>									
Age	2.63	2.71	2.55	2.88	2.43	2.68	2.33	2.70	2.89
sd	(2.09)	(2.11)	(2.07)	(2.07)	(2.01)	(2.17)	(1.82)	(2.19)	(2.27)
Female share	.51	0	1	.19	.44	.81	.51	.17	.58
Dutch share	.75	.76	.73	.83	.67	.77	.75	.89	.72
<i>Level</i>									
Level 2	.26	.31	.22	.32	.26	.21	.27	.35	.24
Level 3	.20	.21	.19	.15	.24	.19	.19	.24	.20
Level 4	.54	.49	.60	.53	.50	.60	.54	.41	.57
<i>Field</i>									
STEM	.24	.40	.09				.21	.68	.18
Economics	.40	.46	.35				.50	.23	.35
Health	.36	.14	.57				.29	.10	.47
<i>Track</i>									
Class-based (BOL)	.81	.75	.87	.67	.82	.90	.82	.58	.85
Apprentice (BBL)	.19	.25	.13	.33	.18	.10	.18	.42	.15
Observations	322,205	156,988	165,217	76,619	129,972	115,614	143,404	28,430	150,371
Share	1	.49	.51	.24	.40	.36	.45	.09	.47

Source: Authors' calculations using non-public microdata from CBS. LSS/HSS stands for low/high skill services.

school level.

For the entire sample, the top 10 verb noun combinations contain three social skills, related to having conversations, either in Dutch, English or a different language. There are three skills related to working with certain (ICT) applications or machines, the remaining four verb noun combinations relate to interpreting specifications, making analyses, writing in English and applying principles.

We see variation in the returns to verb noun combinations between the two tracks in the mbo: class-based (BOL) versus apprenticeship-based (BBL). In the latter, students learn most of their skills on the job, whereas students in the class-based track spend more time learning skills in the classroom. Interestingly, the apprenticeship-based track sees higher returns for learning social skills, such as discussing with colleagues, sales techniques and giving feedback, whereas the top 10 verb noun combinations in the class-based track are less socially oriented.

The pattern for these returns also differs strongly across STEM, economics and health programs. The most obvious differences are in the importance of social and

Table 3.3: Verb Noun Combinations that Positively Correlate Most Strongly with Hourly Wages in one Year after Graduation

Rk	(1) Full sample	(2) Class-based (BOL)	(3) Apprenticeship-based (BBL)
1	Converse English	Contribute	Use ICT application
2	Write English	Apply equipment	Execute calculation
3	Have conversation	Apply machine	Discuss colleague
4	Use equipment	Read English	Apply skill
5	Apply principle	Read assignment	Apply guideline
6	Use language	Apply quality requirement	Apply sales techniques
7	Use ICT application	Apply safety requirement	Understand information
8	Interpret specifications	Converse English	Make calculation
9	Use application	Work system	Give feedback
10	Make analysis	Write English	Apply firm

Rk	(4) STEM	(5) Economics	(6) Health
1	Read information	Read manual	Apply sales techniques
2	Use tool	Apply task	Apply ICT application
3	Act as contact	Apply meeting technq	Apply meeting technq
4	Follow guideline	Apply skill	Apply ICT skill
5	Use equipment	Interpret drawing	Have insight
6	Make calculation	Apply ICT skill	Contribute
7	Keep records	Gather information	Apply presenting technq
8	Make analysis	Have conversation	Apply skill
9	Use application	Write English	Give feedback
10	Convey information	Converse English	Use skill

Rk	(7) Level 2	(8) Level 3	(9) Level 4
1	Read text	Work equipment	Apply equipment
2	Apply rules	Write English	Apply machine
3	Apply task	Apply principle	Contribute
4	Apply skill	Use language	Use ICT application
5	Write English	Have conversation	Have conversation
6	Contribute	Keep records	Converse English
7	Apply information	Use material	Use tool
8	Work equipment	Use application	Interpret drawing
9	Converse English	Converse English	Write English
10	Act as contact	Use equipment	Read information

	Technical
	Social
	Basic

Note: Ranking of coefficients, estimated from a regression on the average hourly wage in the year after graduation on a dummy for whether a verb noun combination is part of the curriculum. A separate regression is estimated for each verb noun combination. Controls are: level, field and track of education, age, gender, and school dummies. The verb noun combinations included in the analysis appear in at least 10 programs, such that positive coefficients are not implicit degree-returns.

technical skills. Especially in health programs, students see higher returns on social skills (highlighted in the darkest shade). Interestingly, these are not necessarily social skills directly related to typical healthcare activities, but more in social skills relating to managing and cooperating in teams: presenting, meeting, sales and negotiation activities. For STEM programs, technical skills (i.e. using certain tools, equipment or applications) are naturally more strongly correlated with wages, but also specific social skills regarding conveying information and taking up a role as a contact in professional settings. Only in STEM programs, math skills are part of the top verbs. For economics programs, the top correlating verb noun combinations are more related to reading skills or processing information. Furthermore, whereas the highest-returning social skills in health curriculums are related to teamwork, we see that the type of social skills that correlate most strongly with economics students are related to having conversations. Even though some skills might be social or technical in nature, different types of each of the three skills are relevant for different fields of education, e.g., social skills relevant for the STEM degree may not be the same as those in, say, health.

Next, columns 5 to 7 show the same results for subsamples of the three levels of middle education. Level 2 students have higher returns to quite general verb noun combinations, related to applying skills, tasks and rules. This would be in line with the fact that level 2 degrees are more related to acquiring basic (learning) skills, rather than specialised technical or social skills. For levels 3 and 4, we see that more specific technical and social skills seem relevant, where especially level 4 students have relatively higher returns to verb noun combinations related to using certain machines, tools or ICT applications. In both cases, both levels 3 and 4 students' have highest returns to social skills related to having conversations, either in Dutch or in another language.

However, using each individual verb noun combination in itself we cannot conclude whether social, technical or basic skills have different returns across subgroups of students. Therefore, in the next section we use the relative skill frequency measures of social, technical, and basic skills to indicate whether some fields or levels of education can benefit more from certain generic skills.

3.5.1 Returns to skill frequencies

To further dive into the types of skills that are rewarded upon entry in the labour market, we estimate a Mincer (1974) equation of wages on the relative frequency of skills in the curriculum. The outcomes show whether curricula that are relatively more focused on certain skills, e.g. social skills, result in better starting wages in the

first year after graduation.

Given the considerations outlined before, our empirical strategy is as follows. Using the skill frequencies as constructed following equation (3.1), we estimate the following regression:

$$\ln w_{ijt+1} = \beta_0 + \beta_1 \text{social}_j + \beta_2 \text{tech}_j + \beta_3 \text{basic}_j + X_i \gamma + Z_j \alpha + \pi_t + \varepsilon_{ij} \quad (3.2)$$

where X_i contains a vector of the demographic controls gender, migration background (first or second generation migrant, and place of origin¹⁶), and age at graduation. Z_j is a vector containing degree-related controls: field of education (STEM, economics and health), level of education (2, 3 or 4), and track (class-based or apprenticeship-based). We include time dummies π_t as the year of graduation.

Besides a baseline regression, we also add interactions between the three skills and i) the levels of education, ii) fields of education and iii) class-based (BOL) versus apprenticeship-based (BBL) tracks. The results in Table 3.4 show how skill-demand differs for students from varying schooling backgrounds. Note again that the sample of students is restricted to those that decided to move to the labour market after graduation from middle education. Students that continued education or have obtained a degree in higher-education following their middle education degree are excluded from the analysis. As such, the results show the returns to skills, conditional on the decision of students to not continue studying.

The results show the following pattern. First, the returns to social-skill intensive curricula are negative across all estimations - even when including interactions. Technical skills are not significantly positive in the baseline estimation, but become positive and significant in all three estimations that include interaction terms. This highlights that especially technical skills are not necessarily valued in each and every curriculum, but that there are level- and field-specific effects at work. Basic cognitive skills are only significant once interactions with levels of education are included in the model.

Next, the bottom panel of the table shows the interactions between levels (column 2), fields (column 3) and tracks of education (column 4). We find little evidence for strict linearity in the relation between skills and levels of education. Social skills are relatively valued more in level 3 programs than the reference category level 2, but are not significantly different from level 4 students. Technical and basic skills are valued

¹⁶This is a variable taking seven options for the most common (migration) backgrounds. These countries, and there approximate share in the total population of mbo-students in 2015, are: Netherlands (73%), Turkey (4%), Morocco (4%), Suriname (4%), Dutch Caribbean (2%), Western (6%), and other non-Western (6%). Source: CBS Statline.

Table 3.4: OLS Regressions of Log Hourly Wage in First Year after Graduation on Skill Frequencies in Curriculum and Interactions with Degree Characteristics

	(1)	(2)	(3)	(4)
Social	-.019*** (.002)	-.024*** (.003)	-.017*** (.005)	-.020*** (.002)
Tech	.016*** (.004)	.008* (.004)	-.012** (.005)	.014*** (.004)
Basic	-.002 (.003)	.004 (.004)	-.013** (.006)	-.000 (.003)
<i>Ref. cat.: Level 2</i>				
Level 3	.108*** (.006)	.106*** (.006)	.076*** (.006)	.108*** (.006)
Level 4	.238*** (.006)	.238*** (.005)	.189*** (.006)	.237*** (.006)
<i>Ref. cat.: STEM</i>				
Econ	-.045*** (.005)	-.043*** (.005)	-.069*** (.005)	-.048*** (.004)
Health	.075*** (.008)	.075*** (.008)	.113*** (.012)	.072*** (.007)
Female	-.035*** (.003)	-.035*** (.003)	-.042*** (.003)	-.036*** (.003)
BBL	.182*** (.008)	.181*** (.008)	.167*** (.008)	.176*** (.008)
Skill frequencies interacted with				
	Level	Field	Track	
	<i>Ref.: Lvl 2</i>	<i>Ref.: STEM</i>	<i>Ref.: BOL</i>	
	Social ×	Social ×	Social ×	.015***
	Level 3	Econ	BBL	(.006)
	Social ×	Social ×		
	Level 4	Health		
	Tech ×	Tech ×	Tech ×	.012**
	Level 3	Econ	BBL	(.006)
	Tech ×	Tech ×		
	Level 4	Health		
	Basic ×	Basic ×	Basic ×	-.016**
	Level 3	Econ	BBL	(.007)
	Basic ×	Basic ×		
	Level 4	Health		
Constant	.770*** (.027)	.773*** (.027)	.848*** (.025)	.781*** (.028)
Obs	322,205	322,205	322,205	322,205
R-squared	.233	.234	.243	.234

Note: *** p<.01, ** p<.05, * p<.1. All models are weighted with robust standard errors, clustered at the school level. Each regression includes the full set of demographic, degree, year, and school controls. BOL stands for class-based track, BBL for the apprenticeship-based track.

most in level 2 students, as indicated by the negative coefficients for level 3 and 4 for both skills. Basic skills have lower returns in level 4, than level 3. Each of these three types of skill thus interact differently with varying levels of education. These results can either hint towards the existence of ability-biased skill returns: some types of skills are mostly valued in combination with certain levels of education. Alternatively, returns to some skills might be higher if they are relatively scarce in degrees of a certain level. This might explain why technical skills have the highest returns for level 2 graduates, whose curricula are relatively more focused towards foundational skills.

Column (3) shows the results with the inclusion of field-skill interactions. Social skills have no field-specific returns. Technical skills are valued more in STEM than in economics programs, and technical skills seem to be especially valued more for health care students, indicated by the coefficient of 0.14. In this case, increasing technical skills in a curriculum thus works best for programs in the health care sector, whereas economics students are more benefited by increasing basic cognitive skills as part of their curriculum.

Lastly, column (4) shows the interactions with tracks: students who follow the apprenticeship track (BBL) have higher returns to social and technical skills than those in the class-based track (BOL). The increase in wages for BBL-students in social skills even offsets the negative general coefficient of -0.018, implying that social and technical skills both significantly relate to higher wages for students in this track. On the other hand, basic skills have lower returns in students of the apprenticeship track. Therefore, the manner in which skills are taught apparently influences the size and sign of the returns to certain types of skills. Social and technical skills relate to higher returns if they are learned by spending relatively more time practicing these skills on the job, whereas basic cognitive skills result in higher returns when learned in class.

3.5.2 Decomposing by sector of employment

So far, we have estimated the returns to skill based on the schooling background of each student: their field, the skills required and the level of education. In this section we decompose the results by sector of employment. We re-estimate equation (3.2) conditional on being employed in either the low skilled services, goods or high skilled service sector. We control for the field and level of education, such that the results obtained here should not capture matches between the sector of education and the sector of employment. The results are presented in Table 3.5.

Intuitively, social-skill intensive curricula should result in high returns in service

Table 3.5: OLS Regressions on Skill Frequencies in Curriculum, by Sector of Employment

	(1) Low-skilled services	(2) Goods	(3) High-skilled services
Social	-.000 (.001)	-.010* (.005)	-.035*** (.002)
Tech	.003 (.002)	-.025*** (.005)	.031*** (.006)
Basic	.001 (.002)	-.018** (.008)	.000 (.003)
<i>Ref. cat.: Level 2</i>			
Level 3	.077*** (.005)	.119*** (.007)	.107*** (.009)
Level 4	.198*** (.005)	.249*** (.010)	.246*** (.008)
<i>Ref. cat.: STEM</i>			
Services	-.047*** (.004)	-.130*** (.015)	-.001 (.007)
Health	-.008 (.005)	-.061*** (.019)	.163*** (.013)
<i>Ref. cat.: BOL</i>			
BBL	.116*** (.006)	.273*** (.009)	.227*** (.012)
Female	-.024*** (.002)	-.066*** (.008)	-.048*** (.004)
Constant	.502*** (.034)	.755*** (.023)	1.060*** (.031)
Observations	143,404	28,430	150,371
R-squared	.260	.341	.212

Note: *** $p < .01$, ** $p < .05$, * $p < .1$. All models are weighted with robust standard errors, clustered at the school level. Each regression includes demographic, degree (track, level and field), year, and school controls, and a constant. BOL stands for class-based track, BBL for the apprenticeship-based track.

sectors - that generally use social skills more intensively. However, this seems not to be the case for middle-educated graduates. The results in Table 3.5 show that socially-oriented curricula are associated with lower returns in the high skill services sectors, and insignificant in the low-skilled services sector. In the goods sector, the return to social skills in curricula is also negative, though it is slightly smaller and less significant.

On the other hand, the returns to technical skills are positive in the high-skilled services sector, yet negative in the goods sector. This is an interesting result, and could possibly be explained by constrained supplies of technical oriented students in the high-skilled services sector, whereas these skills are complementary to the tasks executed in this sector. Regardless, there seem to be sector-specific factors that influence skill-based pay differentials.

3.5.3 Robustness

We perform a few robustness tests to our baseline results from Table 3.4. First, we test the sensitivity of our results to the in- and exclusion of degree-specific controls. The estimations are presented in Table 3.6. First, we would like to test whether our results remain consistent over the inclusion and exclusion of level, field and sub-field controls. Column (1) shows the baseline estimation, as also presented in Table 3.4. In Column (2), field controls are removed, and in column (3) level controls are removed. The results on social skills are highly robust to the exclusion of these variables. The estimates of technical skills also remain positive and significant, but appear to be more sensitive to the exclusion of field and level controls. Next, in column (4) we include sub-field controls, beyond field controls. These sub-fields (or domains) are also presented in Figure 3.3. We see that the effect sizes of both technical and social skill increase in size. Our results are thus robust to the inclusion of more fine-grained field controls. Social (technical) skills are thus robustly and positively (negatively) associated with wages in the first year after graduation, even within disaggregated fields of education.

Lastly, we test whether our results are robust to the sample of verb noun combinations used to construct the skill frequency measure. In the baseline estimations, we use the constructed measures based on 482 verb noun combinations. This sample was selected on the condition that they should appear in at least 5 programs. We recreate this measure based on 152 verb noun combinations that appear in at least 10 programs. To check whether constructing the skill measure based on more common words result in different estimates, we rerun the baseline equation. The results are presented in Column (5). We can see that the results for social skills are robust for

Table 3.6: OLS Regressions on Skill Frequencies in Curriculum, by Sector of Employment

	(1) Base- line	(2) No field controls	(3) No level controls	(4) Add domain	(5) Skill sample
Social	-.019*** (.002)	-.013*** (.002)	-.011*** (.003)	-.024*** (.002)	-.016*** (.002)
Tech	.016*** (.004)	.008*** (.003)	.040*** (.005)	.034*** (.005)	-.001 (.002)
Basic	-.002 (.003)	-.017*** (.002)	.048*** (.003)	.005 (.004)	-.002 (.002)
Level 3	.108*** (.006)	.105*** (.006)		.123*** (.005)	.110*** (.006)
Level 4	.238*** (.006)	.246*** (.006)		.247*** (.006)	.235*** (.006)
Services	-.045*** (.005)		-.009 (.008)	.186*** (.021)	-.058*** (.005)
Health	.075*** (.008)		.140*** (.015)	.090*** (.021)	.054*** (.005)
Female	-.035*** (.003)	-.007** (.004)	-.029*** (.003)	-.030*** (.003)	-.033*** (.003)
Constant	.770*** (.027)	.704*** (.029)	.682*** (.034)	.811*** (.034)	.775*** (.027)
Observations	322,205	322,205	322,205	322,205	322,205
R-squared	.233	.222	.194	.244	.232

Note: *** p<.01, ** p<.05, * p<.1. All models are weighted with robust standard errors, clustered at the school level. Each regression includes demographic, degree (track, level and field), year, and school controls.

the sample size. Technical skills are insignificant, which is most likely due to the fact that specific technologies are mentioned in fewer programs. Therefore, the positive estimate for technical skills in our baseline estimates are likely to be lower-bound estimates, given that they could include even more technology-specific skills if we would include more verb noun combinations in our sample. Basic cognitive skills remain insignificant.

Skill returns in several years after graduation

Since the results in the baseline estimations are all for only one year after graduation, we also run analyses for wages later in the careers of the students: 2, 4, 6, 8 and 10 years after graduation. The results are presented in Appendix Table 3.B.2. The coefficients of the full sample, as well as split out by sector of employment, are graphically presented in Figure 3.4.

We find that our results are not only an artifact of the first year after graduation, but remain persistent in the following years as well. The coefficient of social skills

Figure 3.4: Coefficients of Returns to Social and Technical Skills in Multiple Years after Graduation

Note: Coefficients of regressions presented in Table 3.B.2 with full set of controls. Insignificant estimates are valued at 0. LSS/HSS stand for low skill service sector and high skill service sector respectively.

remains significantly negative over time. It is most strongly negative in the high skill services sector, but this effect decreases over time. The same holds for the positive effect of technical skills, this also decreases in importance over time. This is sensible, as it is likely that the impact of skills taught in school reduce over the period after school, where learning on the job becomes increasingly relevant.

Skill combinations

Existing literature (see, for instance, Deming (2017) and Deming and Kahn (2017)) highlights the fact that combinations between skills explain pay differentials between workers. To test whether we can find similar evidence, we rerun equation (3.2) with the inclusion of interactions between the three skills. We perform this exercise for the full sample, but also for subsamples by sector of employment. The results are presented in Appendix Table 3.B.1.

We find little evidence for skill complementarities in curricula: there are no positive interactions between any of the three skills. This does not need to imply that these skill complementarities do not exist for middle educated workers. However, we find no evidence that including both in a curriculum positively affects starting wages.

Returns to skill factors

The results on social, technical and basic skills highlight that different fields, levels and tracks result in different skill returns. However, the data allow us to go one level deeper, to find out which types of social, technical or basic skills are valued more in these different fields. To do this, we describe the results that use the 36 factors from our factor model.

We estimate the following regression:

$$\ln w_{ijt+1} = \beta_0 + \sum_{f=1}^F \beta_f f_j + X_i \gamma Z_j \alpha + \pi_t + \varepsilon_{ij} \quad (3.3)$$

For F containing 36 factors f on the degree level. The results are presented in Table 3.7. Column (1) shows the baseline estimation of (3.3), columns (2) to (4) are estimations for subsamples of graduates employed in the low-skilled services, goods or high-skilled services sector. The factors are split up by the type of skill: technical, social and basic skills. Below, we only discuss positive correlations. As each skill factor might have an intrinsic value in and of itself, we refrain from the discussion of negative factors in this section. Note that all factors are standardised across the entire sample. For a list of the descriptions of these factors, see Table 3.A.1.

Technical factors that positively correlate with wages are Equipment Maintenance (f5), Management of Maintenance Equipment (f6), Operation and Control (f7), Installation (f10), Following Technical Regulations (f17) and Using Applications (f34). Social factors with positive returns are Coordination in Firm (f1), Management of Human Resources (f4), Conversation Techniques (f26). For basic skills, only Interpreting Manuals (f21) and Following Regulations (f25) are significant and positive.

The results presented here hint towards variation in sector-skill complementarities, and call for further research into understanding sector-specific skill demand. A broad perspective on social skills would overlook the fact that some social skills are more associated with high wages in the high and low skill sector (e.g. Conversation techniques (f26)), whereas others have positive wage returns in the goods sector (e.g. Coordination in Firm (f1)).

A short note on skills versus competencies

The skills as described in this paper concern the specific learned abilities that students will need to perform a job their degree prepares for. A broader view on a person's abilities would also take into account the knowledge and attitudes a person has,

Table 3.7: OLS Regression of Log Hourly Wage in First Year after Graduation on 36 Factors of Verb Noun Combinations

		<i>f</i>	(1) Full	(5) LSS	(6) Goods	(7) HSS
Technical	Mgmt of Material Resources	2	-.020*	.011	-.002	-.052***
	Quality Control Analysis	3	-.003*	-.004	-.007**	.016***
	Equipment Maintenance	5	.034***	.028***	.016***	.041***
	Mgmt of Maintenance Equip	6	.050***	.040***	.021***	.078***
	Operation and Control in Manuf	7	.022***	.002	.011*	.013
	Installation	10	.009**	.012***	.017***	.011
	System Quality Analysis	11	-.001	-.013	-.013	.047**
	Installation in Construction	15	-.014**	.004	.016**	-.050***
	Following technical regulation	17	.008***	.002	-.005	.011***
	Apply Quality Procedures	22	-.010**	-.004	-.015**	-.014***
	Coordination in Operations	23	.015*	-.013	-.001	.012
	Using Tools	27	-.017***	-.006	.004	-.037***
	Tool Selection	29	-.009***	-.010***	-.017*	-.010**
	Reading Equipment Instructions	31	-.010***	-.003	.003	-.013***
	Using Applications	34	.006**	.004	-.008	-.002
	Using Systems	35	-.010***	-.001	-.017**	-.018***
Social	Coordination in Firm	1	.015***	-.003	.017***	-.006
	Mgmt of Human Resources	4	.030***	.008***	.006*	.033***
	Mgmt in Office setting	13	-.006***	-.005***	-.013***	-.004*
	Conservation techniques	26	.010***	.006***	.002	.016***
	Persuasion	28	-.006	-.010***	-.001	-.006
	Sales	30	-.007***	-.004***	-.010***	-.005***
	Administrative Work	32	-.002	-.007	-.001	.014*
	Discussing Calculations	33	.002	.002	.003	.007***
Basic	Reading/Following Instructions	8	.003	.005*	-.003	.007
	Discussing Size Calculations	9	-.007***	-.002	-.011***	-.010***
	Reading in ICT setting	12	-.021**	-.011*	.013	-.038***
	Apply skill	14	-.016***	.004	-.022***	-.039***
	Field-specific Reading	16	.001	.006***	.001	-.002
	Complex problem solving	18	-.013***	-.010***	-.026***	-.021***
	Foreign language	19	-.005***	-.005**	-.015***	-.005**
	Reading Specifications	20	-.002	-.010***	-.000	.010
	Interpreting Manuals	21	.010***	.009***	-.005*	.017***
	Reading Comprehension	24	-.009***	-.011***	-.005	-.008
	Following regulation	25	-.001	.004**	.008	-.005
	Following guidelines	36	-.003***	.004***	.003	-.009***
Controls	Level 3		.085***	.087***	.115***	.063***
	Level 4		.240***	.221***	.282***	.243***
	Economics		-.064***	-.043***	-.114***	-.057***
	Health		-.023**	.000	-.015	-.022
	Female		-.042***	-.020***	-.076***	-.057***
	Constant		.846***	.484***	.727***	1.139***
	N		322,205	143,404	28,430	150,371
	R ²		.254	.267	.359	.233

Note: *** p<.01, ** p<.05, * p<.1. All models are weighted with robust standard errors, clustered at the school level (not reported here). Each regression includes demographic, degree (track, level and field), year, and school controls, and a constant. LSS/HSS stands for low/high skill service sector.

which will determine their capacity to successfully execute a task. Focusing on skills alone overlooks this, and bears the risk of simplifying a worker to the tasks they execute - rather than the attitudes and knowledge in the background that define the success of a task. Competencies comprise the combination of skills, knowledge and attitudes. As a term, it is receiving increasing attention in the education sector. Examples of such competencies are flexibility, resilience, ambition, decision-making, or networking.

Besides skills, each curriculum contains a list of competencies (see Figure 3.2). Appendix 3.C shows tentative results of the returns to competencies in the curriculum. The curriculum-data on competencies is far from perfect: even more than for the skills data, curriculum writers may have varying interpretations of competencies. Nevertheless, the appendix shows that these competencies might be an interesting future avenue of research - worth adding to the growing literature on skills.

3.6 Conclusion

This chapter uses novel data on skills to estimate returns to specific elements in the curricula of Dutch middle-educated students. We created skill measures based on the relative frequency of social, technical and basic cognitive skills mentioned in the curriculum. We performed two main exercises with this data. First, we showed how skill supply differs between sectors and levels of education, in terms of what is being taught in the curricula. Second, we linked the curriculum contents to wage data in the first years after graduation, to estimated returns to granulated skill levels.

We found that students graduating from degrees with a stronger focus on social skills have lower wage outcomes in the first years after graduation, even after controlling for specific fields of education. A decomposition analysis by sector shows that this seems to be mainly driven by returns in the high skill services sector. Given the nature of our data, we believe our estimates have little bias, since students select into degrees, not into verb noun combinations, and since the curricula are constructed at a national level. This makes our skill measure of observation relatively exogenous to the student choice.

Furthermore, we would like to stress that there is no ‘one-skill-fits-all’ for Dutch middle educated students. The returns to social, technical and basic skills strongly differ across fields of education, sector of employment, level of education, and class-based or apprenticeship based tracks. We also showed that learning social skills in school may not be beneficial for middle-educated students upon entering the labour market. Nevertheless, it might be the case that social skills will pay off later

in their careers - if they are learned on the job. The results from the factor analysis and the regression on factors also highlight that social, technical and basic skills are multifaceted terms and different aspects of each of those skills are needed in different sectors.

Importantly, our findings imply that those who enroll in a degree with a relatively stronger focus on obtaining social skills have lower wages than degrees focusing on other types of skills, e.g., technical skills. The methodology employed and the results provided could be used for further research on the topic. For instance, it could be further explored that rather than encouraging students to enter a STEM degree, it might be a better policy to increase the relative importance of technical skills in all degrees - and especially for those students who are educated in the field of healthcare, or working in the high skill services sector. Overall, our results address some of the key aspects of middle education in the Netherlands, with implications for a number of other countries using similar education systems, and its labour market implications.

3.A Additional descriptive tables and figures

Figure 3.A.1: Screeplots of Eigenvalues for factors, on more and less restrictive sample

(a) Sample: 152 verbs, each verb in at least 10 programs. 36 factors with Eigenvalue > 1.

(b) Sample: 482 verbs, each verb in at least 5 programs. 80 factors with Eigenvalue > 1.

Table 3.A.1: Factor Descriptions, sorted by type of Skill following the O*NET classification of Job Skills

<i>f</i>	Skill type	Factor description	Factor elements
2	Technical skills	Management of Material Resources	Meet, save and maintain materials, products, provisions, tools
3		Quality Control Analysis	Work quality demands, using testing equipment/tools, looking up instructions/systems
5		Equipment Maintenance	Maintain machines/tools, recognize/process materials, apply quality norm
6		Management of Maintenance Equipment	Maintain/use protective equipment, maintain equipment
7		Operation and Control in Manufacturing	Operate means of transport, shield waste container
10		Installation	Edit component/part/material
11		System Quality Analysis	Apply systems, check quality, apply machines
15		Installation in Construction	Working with blueprint, slinging loads
17		Following technical regulation	Apply environmental rule/company regulation
22		Apply Quality Procedures	Apply quality demand, apply procedure
23		Coordination in Operations	Operate machine, use protective equipment, cooperate with colleague
27		Using Tools	Apply tool/protective equipment/materials
29		Tool Selection	Selecting material/tool, work with equipment
31		Reading Equipment Instructions	Interpret scheme, execute task, use equipment
34		Using Applications	Apply part, use application
35		Using Systems	Work system
1	Social skills	Coordination	Distinguish task/function/responsibility, report in and contribute to meetings
4		Management of Personnel Resources	Recognise problem, act as contact, feedback/reflection/observation skills, instructing others
13		Management in Office setting	Apply ICT skills/meeting skills/feedback skills/presenting skills
26		Conversation techniques	Apply knowledge/tools/conversation techniques
28		Persuasion	Convey information, apply negotiation technique, make decision
30		Sales	Apply presenting skills/sales techniques/techniques
32		Administrative Work	Keep up administration, apply observation techniques, take action, use registration system
33		Discussing Calculations	Perform calculation, give feedback, cooperate colleague
8		Reading and Following Instructions	Read/use checklist/assignment, consult source of information
9		Discussing Size Calculations	Calculate dimensions, discuss with colleague
12		Reading in ICT setting	Understand/read assignment, use ICT applications/systems
14		Apply skill	Apply skill
16		Field-specific Reading	Reading field, reading text
18		Complex problem solving	Prepare SWOT analysis, make analysis, interpret data, make calculation
19		Foreign language	Converse/write in English, converse in language
20	Basic skills	Reading Specifications	Reading/interpreting specifications and blueprints
21		Interpreting Manuals	Reading manuals, interpret schemes, have insight
24		Reading Comprehension	Read documentation/information/instruction/English, make drawing
25		Following regulation	Apply regulation
36		Following guidelines	Work guideline

Table 3.A.2: Classification of industries into sectors, employment shares and skill frequencies

Sector	SBI	Industry	Employment				Shares (in %)				Skill freq (std)			
			Total	STEM	Econ	Health	STEM	Econ	Health	Social	Tech	Social	Tech	Basic
Low-skilled services	G	Wholesale and retail trade	83,745	18,541	42,546	22,658	22	51	27	.65	-.23	.65	-.23	-.30
	H	Transportation and storage	8,139	3,875	3,183	1,081	48	39	13	0.41	-0.19	0.41	-0.19	-0.31
	I	Accommodation and food services	37,369	5,277	22,194	9,898	14	59	26	0.76	-0.31	0.76	-0.31	-0.51
	R	Culture, sports and recreation	7,517	1,150	2,436	3,931	15	32	52	1.14	-0.40	1.14	-0.40	-0.58
	S	Other service activities	6,634	1,477	1,125	4,032	22	17	61	-0.15	-0.56	-0.15	-0.56	-0.62
		Subtotal	143,404	30,320	71,484	41,600	21	50	29					
Goods	A	Agriculture, forestry and fishing	3,330	1,309	1,170	851	39	35	26	0.54	-0.20	0.54	-0.20	-0.24
	B	Mining and quarrying	55	n.d.	n.d.	n.d.	n.d.	n.d.	n.d.	-0.40	-0.19	-0.40	-0.19	0.55
	C	Manufacturing	14,053	8,545	3,936	1,572	61	28	11	-0.07	0.18	-0.07	0.18	-0.01
	F	Construction	10,992	9,377	1,284	331	85	12	3	-0.42	0.09	-0.42	0.09	0.18
		Subtotal	28,430	19,231	6,390	2,754	68	22	10					
High-skilled services	D	Electricity, gas, steam and air cond supply	287	212	60	15	74	21	5	-0.24	0.25	-0.24	0.25	-0.33
	E	Water supply	314	198	86	30	63	27	10	0.14	0.06	0.14	0.06	-0.33
	J	Information and communication	5,559	1,182	3,748	629	21	67	11	0.76	-0.40	0.76	-0.40	-0.51
	K	Financial institutions	1,981	344	1,158	479	17	58	24	0.63	-0.35	0.63	-0.35	-0.24
	L	Renting, buying and selling of real estate	1,084	211	594	279	19	55	26	0.63	-0.34	0.63	-0.34	-0.20
	M	Professional, scientific and technical services	11,267	3,480	5,973	1,814	31	53	16	0.60	-0.18	0.60	-0.18	-0.18
	N	Renting and leasing, and support services	62653	15975	31031	15,647	25	50	25	0.59	-0.28	0.59	-0.28	-0.33
	O	Public administration	4,767	1,140	2,827	800	24	59	17	0.13	-0.40	0.13	-0.40	-0.28
	P	Education	6,822	2,631	1,441	2,750	39	21	40	0.39	-0.18	0.39	-0.18	-0.41
	Q	Human health and social work	55,637	1,655	5,166	48,816	3	9	88	1.23	-0.36	1.23	-0.36	-0.85
		Subtotal	150,371	27,028	52,084	71,259	18	35	47					
		Total	322,205	76,579	129,958	115,613	24	40	36					

Note: Classification of industries into sectors based on Bárány and Siegel (2020). Last three columns represent standardised values of skill frequencies.

Figure 3.A.2: Standardised factor scores by major sector

Figure 3.A.3: Standardised factor scores of selected factors, by domain

Source: Authors' calculations using non-public microdata from Netherlands statistics. Note: The value of factor 26 for 'Catering and Bakery' is capped at from 2.8 to 1.5 to accommodate common scales. The domain 'Safety and Sports' contains some economics programs as well as health programs, which is why it has two bars.

3.B Robustness Tables

Table 3.B.1: OLS Regressions of Log Hourly Wage in First Year after Graduation on Skill Frequencies in Curriculum and Interactions with other Skills, by sector and track

	(1) Full	(2) LSS	(3) Goods	(4) HSS
Social	-.031*** (.003)	-.014*** (.003)	-.030*** (.009)	-.039*** (.004)
Tech	.007** (.004)	.005 (.003)	-.043*** (.006)	.022*** (.006)
Basic	-.018*** (.005)	-.003 (.003)	-.046*** (.011)	-.012* (.007)
Social × Tech	-.007** (.003)	-.007*** (.002)	-.031*** (.006)	-.005 (.004)
Social × Basic	-.009*** (.002)	-.012*** (.002)	-.014* (.008)	-.001 (.003)
Tech × Basic	-.016** (.006)	.005 (.004)	-.034*** (.009)	-.017* (.009)
<i>Ref. cat.: Level 2</i>				
Level 3	.113*** (.006)	.083*** (.005)	.131*** (.007)	.107*** (.011)
Level 4	.241*** (.006)	.206*** (.005)	.255*** (.012)	.245*** (.007)
<i>Ref. cat.: STEM</i>				
Service	-.043*** (.004)	-.038*** (.004)	-.122*** (.014)	-.002 (.007)
Health	.073*** (.008)	-.009 (.005)	-.077*** (.022)	.161*** (.012)
<i>Ref. cat.: BBL</i>				
BBL	.184*** (.008)	.116*** (.005)	.275*** (.009)	.227*** (.012)
Female	-.035*** (.003)	-.024*** (.002)	-.067*** (.008)	-.049*** (.004)
Constant	.761*** (.027)	.493*** (.033)	.734*** (.024)	1.056*** (.032)
Observations	322,205	143,404	28,430	150,371
R-squared	.234	.261	.343	.213

Note: ***p<.01, **p<.05, *p<.1. All models are weighted with robust standard errors, clustered at the school level. Each regression includes demographic, degree (track, level and field), year, and school controls, and a constant. LSS/HSS stands for low/high skill service sector. BOL stands for class-based training, BBL for apprenticeship-based training.

Table 3.B.2: OLS Regressions of Log Hourly Wage in 2 to 10 Years after Graduation on Skills

A. Full Sample	(1) t+2	(2) t+4	(3) t+6	(4) t+8	(5) t+10
Social	-.020*** (.002)	-.025*** (.002)	-.020*** (.002)	-.018*** (.002)	-.018*** (.002)
Tech	.021*** (.003)	.016*** (.003)	.011*** (.003)	.011*** (.003)	.010*** (.004)
Basic	-.006*** (.002)	-.010*** (.002)	-.010*** (.002)	-.011*** (.002)	-.011*** (.003)
Observations	320,912	312,708	270,711	199,730	137,528
R-squared	.142	.107	.122	.150	.160

B. Low-skilled services	(1) t+2	(2) t+4	(3) t+6	(4) t+8	(5) t+10
Social	.000 (.002)	-.003** (.001)	-.005*** (.001)	-.006*** (.001)	-.007*** (.002)
Tech	.008*** (.003)	.005* (.003)	.003 (.003)	.003 (.004)	.005 (.004)
Basic	-.000 (.003)	-.001 (.002)	-.006** (.002)	-.013*** (.003)	-.017*** (.004)
Observations	128,317	104,128	79,914	55,521	36,639
R-squared	.157	.076	.093	.141	.170

C. Goods	(1) t+2	(2) t+4	(3) t+6	(4) t+8	(5) t+10
Social	-.017*** (.004)	-.028*** (.004)	-.011*** (.003)	-.014*** (.003)	-.013*** (.003)
Tech	-.018*** (.005)	-.017*** (.005)	-.006 (.004)	.005 (.003)	.010*** (.003)
Basic	-.028*** (.008)	-.021*** (.006)	-.012** (.005)	-.004 (.003)	-.002 (.004)
Observations	31,109	34,819	33,380	26,569	19,576
R-squared	.223	.155	.142	.172	.203

D. High-skilled services	(1) t+2	(2) t+4	(3) t+6	(4) t+8	(5) t+10
Social	-.038*** (.003)	-.040*** (.002)	-.026*** (.002)	-.021*** (.002)	-.019*** (.001)
Tech	.038*** (.005)	.029*** (.005)	.015*** (.004)	.011*** (.003)	.007** (.004)
Basic	-.006** (.003)	-.015*** (.003)	-.011*** (.002)	-.011*** (.002)	-.009*** (.003)
Observations	161,486	173,761	157,417	117,640	81,313
R-squared	.138	.124	.133	.152	.153

Note: *** p<.01, ** p<.05, * p<.1. All models are weighted with robust standard errors, clustered at the school level. Each regression includes the full set of demographic, degree (track, level and field), year, and school controls, plus a constant.

3.C Skills versus competencies: a broader view on abilities

The task-replacing nature of technology has increased the focus in economic literature on the complementarity and substitutability between workers and machines. Given that machines are increasingly able to perform certain tasks, the literature has focused on the specific abilities that machines have versus human labour. Much of the focus in this type of research compares the executability of specific tasks between man and machines. This narrows the view on skill-demand: skills are inputs for the execution of tasks.

However, within the education sector there is a growing debate on whether skills are not too narrow as learning goals. Skills are related to the specific, physical activities that a person is skilled in, for instance a nurse that knows how to place an injection or a carpenter that knows how to use certain equipment. However, in order to fully function in a job, a person needs to combine these skills to knowledge, as well as their own attitudes. This has shifted the focus of skills to a broader term of 'competencies', which comprise knowledge, skill and attitudes as a more complete measure of a person's abilities.

In a separate section of each curriculum file, tasks are linked to a list of predefined competencies, that fit into eight larger categories. Each curriculum writer labels the tasks to competencies, using a document with detailed descriptions and examples of tasks that require certain competencies. There are eight main categories, as presented in Table 3.C.1.

3.C.1 Which competencies are students learning?

The relative prevalence of the 25 competencies are presented in Figure 3.C.1 and 3.C.2.

Figure 3.C.1 shows significant differences between the three major sectors of education. In this case, I only include the sample of level 4 students, to show relative differences for the same and highest level of education within the mbo. STEM degrees are characterised by lesser focus on the core competency 'interact, influence and presenting' and relatively more focus on 'analyzing and interpreting'. This shows that STEM degrees are more focused on learning specific skills relating to their field, such as knowing how to use certain materials and applying expertise. Healthcare students are taught more skills related to coaching ("Begeleiden") and doing research, and less than economics students on commercial acting. Economics students are relatively trained more in competencies as presenting, cooperating, and commercial acting.

Table 3.C.1: Competencies with examples of components

Category	Competency	Examples of components
Managing and Decision-making	Decision-making Supervision Coaching	Make decisions, consider risks, show confidence Delegate tasks, give instructions, exercise authority Coach, advise, motivate, and develop others
Support and Cooperate	Social Perceptiveness Cooperating Ethical conduct	Show interest, listen, empathy, support others Consult others, adapt to group, communicate openly Act with integrity, respect differences between people
Interact, Influence and Present	Networking Persuasion Presenting	Build and maintain relationships, mediate in disagreements Make impression, appeal to emotions, negotiate Explain clearly and concisely, use humor, radiate reliability
Analyse and Interpret	Reporting Apply expertise Using resources Analysis	Formulate correctly, attractively and concisely, add structure Apply subject-specific (manual) skills, share expertise Choose appropriate materials, use means effectively Obtain info from data, draw conclusions, make connections
Create and Learn	Research Innovation Learning	Retrieve information, observe from multiple angles Act creatively, search for change, show vision of future Maintain specific knowledge and skills, learn from feedback
Organise	Planning Customer Orientation Delivering Quality Following Instructions	Set goals and priorities, manage time, organise resources Match needs and expectations, check customer satisfaction Formulate and attain quality norms, work systematically Show discipline, work according to procedures
Adapt and Cope	Flexibility Resilience	Adapt to changing circumstances, accept new ideas Perform under pressure, control feelings, stay positive
Efficacy	Ambition Entrepreneurship Professionalism	Accept challenges, show spirit, take responsibility Know market, identify opportunities, help firm grow Show financial awareness, understand firm dynamics

Moreover, variation in terms of levels shows a clear pattern of complex versus simpler competencies. The skills that are relatively taught more in higher levels of mbo are 'analyzing', 'supervision', the core competency 'create and learn', 'commercial acting' and 'entrepreneurial behaviour'. On the other hand, level 2 students focus relatively more in cooperating, using materials and resources, but also strongly focus on following instructions and delivering quality. In Level 4, we see that students receive a broader spectrum of skills, and are thus not likely to gain deeper learning experiences in each competency due to their higher level, but also a broader set of skills than level 2 students.

Figure 3.C.1: Competencies taught in curriculum, by sector of education.

Source: Author's calculation using S-BB data. Only level 4 students are included in this sample.

Figure 3.C.2: Competencies taught in curriculum, by level of education.

Source: Author's calculation using S-BB data.

3.C.2 Returns to competencies

We use the competencies data provided by S-BB, to estimate the following regression:

$$\ln w_{ijt+1} = \beta_0 + \sum_{c=1}^C \beta_c c_j + X_i \gamma Z_j \alpha + \pi_t + \varepsilon_{ij} \quad (3.4)$$

For C containing 25 competencies c as listed by S-BB. The results are presented in Table 3.C.2, where only the positive and significant coefficients are accentuated. Each of the competencies can be placed in a larger core competency, as can be seen in the first column. The following patterns emerge.

STEM graduates have positive returns to: decision-making, supervision, social perceptiveness, networking, customer orientation, delivering quality, following instructions, and ambition and entrepreneurship. Many of these skills are related to functioning properly in an organisation. We would refer to this as ‘inward’ social skills. Decision-making and supervising others are examples of such inward social skills: they reflect the ability to understand how firms operate and to act based on their social rules. In STEM programs these skills are associated with a positive return. In the category of supporting and cooperating skills, social perceptiveness is related higher wage premiums in STEM graduates. This competency is defined by components as being interested, listening, showing empathy and support, and being understanding of others and oneself. Even though this skill is more often included in health curricula, in the case where it does become part of a STEM degree it is associated with a higher wage premium.

Other positive returns to inward-oriented social skill are “delivering quality” and “following instructions”, which fall in the core category of organizing and planning. These are not significant for economic and health programs. More outward-oriented social competencies like those in the category “Interact, Influence and Present” show a converse pattern: networking, persuasion and presenting are not important skills for STEM graduates. Apparently, STEM graduates are mostly valued when they learn skills that allows them to work according to the standard processes in the firm.

Second, we can observe that economic curricula have a somewhat similar social-oriented pattern of skill returns. Economics graduates have positive returns to decision-making, cooperating, presenting, research, resilience and acquiring an entrepreneurial attitude. Entrepreneurship captures a commercial lookout, identifying threats and opportunity on the market of competitors, which, in combination with skills as presenting, decision-making and research, paints a picture where these students are more ‘outward’ oriented than their STEM peers.

Table 3.C.2: OLS Regressions of Mean Log Hourly Wage in One Year after Graduation on Competencies, for Full Sample and Sub-Samples

Category	Competency	By sector				By level		
		(1) Full	(2) STEM	(3) Econ	(4) Health	(5) Level 2	(6) Level 3	(7) Level 4
Managing and Decision-making	Decision-making	.045*** (.0072)	.045*** (.0083)	.067*** (.011)	.00098 (.0069)	.15*** (.038)	.14*** (.013)	.013 (.013)
	Supervision	.041*** (.0049)	.077*** (.0076)	-.037*** (.0093)	.091*** (.013)	.25*** (.051)	-.00023 (.012)	.074*** (.0065)
	Coaching	-.0033 (.0081)	-.029** (.012)	-.025*** (.0069)	.0020 (.010)	-.0084 (.017)	.049* (.029)	-.015 (.011)
Support and Cooperate	Social Sensitivity	.013 (.013)	.12*** (.032)	.060*** (.0069)	-.063** (.030)	-.12*** (.039)	.13*** (.019)	.047* (.028)
	Cooperating	.031*** (.0071)	-.025 (.018)	.067*** (.012)	.0033 (.020)	.13** (.049)	.095*** (.022)	-.056** (.022)
	Ethical conduct	.038*** (.0059)	-.013 (.0095)	.044*** (.0062)	.097*** (.018)	-.12*** (.021)	.055*** (.018)	-.013 (.011)
Interact, Influence and Present	Networking	.044*** (.0053)	.0069 (.017)	-.087*** (.010)	-.021* (.010)	.23** (.096)	.090*** (.011)	.060*** (.0078)
	Persuasion	-.0031 (.0047)	-.057*** (.014)	-.019*** (.0069)		.18** (.075)	-.059*** (.013)	.017** (.0071)
	Presenting	.0028 (.012)	-.094*** (.017)	.065*** (.010)	.10*** (.021)	.023 (.036)	-.057* (.032)	.051*** (.015)
Analyse and Interpret	Reporting	-.015 (.0098)	-.0070 (.012)	-.012 (.0088)	-.11*** (.033)	-.013 (.015)	.11*** (.024)	-.18*** (.051)
	Apply expertise	-.15*** (.024)		-.12*** (.027)				-.14*** (.035)
	Using resources	-.025*** (.0057)	-.080*** (.014)	.0075 (.010)	-.092*** (.024)	-.076* (.039)	.080*** (.025)	-.030*** (.010)
Create and Learn	Analysis	.017** (.0069)	-.056*** (.0069)	-.053*** (.018)	.20*** (.013)	-.10*** (.015)	.030*** (.0092)	.058*** (.011)
	Research	-.030*** (.010)	-.029* (.017)	-.025*** (.0088)	-.066*** (.017)	.045* (.027)	-.091*** (.017)	-.050*** (.012)
	Innovation	.024*** (.0063)	-.100*** (.017)	-.034*** (.0085)	.042*** (.013)		.11*** (.012)	-.016 (.012)
Organise	Learning	.069*** (.0083)	-.039*** (.0081)	-.0021 (.015)	-.049* (.025)	-.11*** (.018)	-.074 (.053)	.0042 (.024)
	Planning	-.0029 (.012)	.027 (.019)	-.017* (.0097)		-.042** (.017)	-.061** (.027)	.034 (.035)
	Customer Orientation	-.074*** (.016)	-.018 (.014)	-.092*** (.0091)		.056*** (.021)	.024 (.016)	-.093*** (.028)
Adapt and Cope	Delivering Quality	.021** (.0083)	.12*** (.028)	-.0088 (.0059)	-.0035 (.031)	.036* (.020)	.023 (.022)	.036* (.019)
	Following Instructions	.047*** (.010)	.051*** (.015)	-.0090 (.018)		-.046 (.049)	.021 (.055)	.035*** (.013)
	Flexibility	-.0064 (.0080)	.018 (.024)	-.020 (.012)	.060*** (.012)	.15** (.060)	.15** (.063)	.0044 (.0073)
Efficacy	Resilience	.031*** (.0065)	.014 (.012)	-.0058 (.0087)	.068*** (.011)	.057*** (.019)	-.082*** (.012)	.039*** (.011)
	Ambition	-.068*** (.0070)	.034*** (.013)	.066*** (.013)	-.12*** (.0078)	.22*** (.036)	.26*** (.031)	-.085*** (.0081)
	Entrepreneurship	-.041** (.017)	.062 (.037)	.015* (.0085)	-.15*** (.022)	-.18*** (.051)	.077*** (.018)	-.069** (.034)
	Professionalism	-.056*** (.0049)	-.020 (.019)	-.048*** (.0099)	.065*** (.0097)	-.11*** (.030)	-.14*** (.029)	-.055*** (.011)

Note: *** p<.01, ** p<.05, * p<.1. All models are weighted with robust standard errors, clustered at the school level. Each regression includes demographic, degree (track, level and field), year, and school controls, and a constant.

Even though health graduates also have high returns to social skills, their labour market returns also show more relevance for cognitive skills and competencies related to adapting to a changing environment. Health students have positive returns to: supervision, coaching, presenting, analysis, flexibility, resilience, and professional behavior. “Analysis” captures skills related to generating information from data, checking data, drawing conclusions and thinking of solutions for problems, which, especially in combination with innovation, can be seen as a cognitive skill, overlapping most strongly with the O*NET skill of complex problem solving.

Lastly, some competencies are valued more in higher levels of education. Naturally, higher levels of education also typically increase the set of competencies and deepen the abilities of students. Some competencies are almost exclusively taught in higher level programs, such as applying expertise and innovation. Positive returns in high level programs are “supervision”, “social perceptiveness”, “interact, influence and present”, “following instructions” and “flexibility”. Following instructions is interesting here, as it probably relates to more routine tasks when part of a level 2 or 3 program, but less so when included in a level 4 program.

Conclusion on competencies versus skills

In taking a broader view on human skills, one can focus on competencies, which bundle skills with knowledge and attitudes. Adding the results on competencies to the analyses on skill frequencies in the main text allows us to paint a broader picture of labour supply (what students learn) and demand (what students earn) in the Netherlands.

The Dutch results on competency-level data show that the organisation of work is different along the three major sectors of education, which leads to varying returns to general competencies. Specifically, STEM graduates earn more when they learn inward-oriented social skills besides their occupation-specific knowledge. Economics curricula focused on both inward as well as outward social skills have higher returns, and specifically if they include the more ‘extroverted’ outward skills that are marketable and can generate output for the firm: commercial orientation, decision-making and presenting. In contrast to STEM and economics profiles, cognitive and adaptive competencies are associated with positive returns for health graduates.

4

Meaningful Work

4.1 Introduction

The concept of meaningful work—activities that individuals view as purposeful and worthwhile — has received relatively little attention in modern economics.¹ While organisational psychologists have long examined the meaning people derive from their jobs (Rosso, Dekas, and Wrzesniewski, 2010), modern economists typically conceive of work as a disutility, i.e. as an unpleasant activity that must be endured as a means to earn an income and finance consumption. Nevertheless, studies relying on self-reported and experimental data have challenged the assumption that only monetary motivations matter in the labour market (Binder, 2016; Bradler, Dur, Neckermann, and Non, 2016; Hamermesh, 2018; Hamilton, 2000; Kosfeld, Neckermann, and Yang, 2017; Preston, 1989; Stern, 2004). In fact, one convincing piece of evidence for the intrinsic value of work is the enormous psychological cost of becoming unemployed, which by far exceeds income losses (Clark, 2001; Kassenboehmer and Haisken-DeNew, 2009; Knabe and Rätzl, 2011a,b; Nikolova and Ayhan, 2019; Winkelmann and Winkelmann, 1998).

Nevertheless, despite the recent attention to non-monetary work-incentives (Lazear, 2018), only two economics papers have called for incorporating work meaningfulness in standard labour supply models (Cassar and Meier, 2018; Spencer, 2015). Meanwhile, the empirical research on meaningful work in organisational psychology has left several knowledge gaps (Lysova, Allan, Dik, Duffy, and Steger, 2019). While important, these studies rely on small non-representative samples, lack a unified definition of work meaningfulness, and use divergent measurement scales that often conflate meaningfulness with other constructs such as calling (Bailey, Lips-Wiersma, Madden, Yeoman, Thompson, and Chalofsky, 2019a; Bailey, Yeoman, Madden, Thompson, and Kerridge, 2019b). This is unfortunate because it limits our understanding of which factors contribute to work meaningfulness and whether it is conducive to behaviours such as increased effort and delayed retirement.

This chapter is based on the paper “What makes work meaningful and why economists should care about it” in *Labour Economics* (Nikolova and Cnossen, 2020). DOI: 10.1016/j.labeco.2020.101847. The authors have benefited from the feedback of Carol Graham, Clemens Hetschko, Jolanda Hessels, David Spencer, Spyridon Stavropoulos, Ruut Veenhoven, Indy Wijngaards, Markus Krecik, Mark Fabian and seminar participants at the Leeds University Business School and the Erasmus Happiness Economics Research Organisation (EHERO), as well as ISQoLS 2019 conference participants in Granada. Furthermore, we would like to thank the Editor of *Labour Economics*, Arthur van Soest, and two anonymous referees for detailed comments and constructive suggestions. The authors would also like to acknowledge copy-editing support from Julia Bloom and the contributions to an earlier version of this paper by Puck Otten. In addition, we are grateful to the UK Data Service for access to the 2005, 2010, and 2015 Working Conditions Surveys. All errors are our own. Personal contribution to this chapter: data editing, idea development, model creation, writing.

¹Throughout this paper, by “meaningful work,” we mean the individual’s own perceptions of being engaged in meaningful work.

This paper closes these knowledge gaps by making a threefold contribution to the literature: first, we are the first to investigate the determinants and consequences of meaningful work using a cross-country nationally representative dataset of workers from 30 European countries in 2005, 2010, and 2015. Conceptually, we rely on self-determination theory (Deci and Ryan, 1985; Ryan and Deci, 2000b), which outlines the conditions leading to motivation and work meaningfulness. Empirically, we construct an index of meaningful work based on survey statements about perceptions of doing useful work and having feelings of “a job well done” (fulfillment) (see Appendix Table 4.A.2 and Figure 4.1). As such, we are the first to quantify the relative importance of job characteristics that enhance or diminish work meaningfulness, which could help inform policies and interventions to promote work meaningfulness. We find that autonomy, competence, and relatedness are about 4.6 times more important for meaningfulness at work than compensation, benefits, career advancement, job insecurity, and working hours. Relatedness, which reflects supportive relationships with colleagues and superiors, emerges as the most important factor for work meaningfulness. These findings highlight the greater salience of self-efficacy and intrinsic motivation for meaningfulness compared to objective working conditions and monetary rewards.

Second, we show that perceptions of meaningful work have implications for labour economics because they meaningfully predict retirement intentions, absenteeism, and skills training. For example, a ten-point increase in work meaningfulness reduces absenteeism by about one day per year and raises the intended retirement age by 2.5 years, on average. These findings not only validate the usefulness of work meaningfulness in economics but also have relevant implications for employers and policy-makers. Our third contribution is that we outline the conceptual and methodological steps that can contribute to a future research agenda in meaningful work in economics.

4.2 Conceptual Framework

4.2.1 Worker well-being and meaningful work

While well-being is a latent construct, there are two main approaches to conceptualising and measuring it, which have their own advantages and disadvantages (Brown, Charlwood, Forde, and Spencer, 2007; Brown, Charlwood, and Spencer, 2012; Green, 2006; Knox, Warhurst, Nickson, and Dutton, 2015). According to the *objective* approach, worker well-being is about having the capabilities and freedoms that allow individuals to meet specific needs, such as autonomy and skills development (Brown

et al., 2007, 2012; Green, 2006). This framework draws on Sen's (1999) capability approach, which conceptualises well-being in terms of having the capabilities and freedoms to achieve 'functionings', i.e. states of being and doing that the individual values. In the objective approach, worker well-being can be evaluated based on whether the job furnishes workers with the capabilities and material security to achieve their goals and fulfill their needs (Brown et al., 2012; Budd and Spencer, 2015; Eurofound, 2012; Green, 2006). As such, the objective approach relies on 'job quality' measures, i.e. job characteristics and working conditions (De Bustillo, Fernández-Macías, Antón, and Esteve, 2011b; Felstead, Gallie, Green, and Henseke, 2019; Green, 2006; Howell and Kalleberg, 2019). For example, based on Sen's capability approach, Green (2006) defines job quality in terms of earnings, skill, effort, autonomy, security, and personal discretion. The extent to which a person's job provides these factors determines their ability to achieve well-being (Budd and Spencer, 2015). To date, several multi-dimensional job quality indices have been created (De Bustillo et al., 2011b; Eurofound, 2012; Leschke and Watt, 2014). Nevertheless, the main challenges of the objective approach concern the choice and measurement of the characteristics that comprise job quality (Budd and Spencer, 2015; Clark, 2011). In particular, the final list of selected job quality measures may reflect data availability and researcher discretion, rather than worker preferences. Challenges may also arise when picking what weights the job quality measures should receive to form a comprehensive multi-dimensional job quality index (Leschke and Watt, 2014; De Bustillo, Fernández-Macías, Esteve, and Antón, 2011a). Because it focuses on job characteristics, the objective approach has also been criticised for being job-centric and for ignoring the broader meanings that work has in people's lives (Budd and Spencer, 2015).

In contrast, the subjective well-being approach assumes that people themselves are the best judges of their working and living environments (Graham, Laffan, and Pinto, 2018; Graham and Nikolova, 2015; MacKerron, 2012; OECD, 2013; Stone and Mackie, 2014). In the work domain, subjective measures of well-being include self-reported feelings and evaluations of the overall working conditions. In economics, the most common subjective well-being measure in the work domain is job satisfaction (Clark, 2001, 2005, 2011, 2015; Sousa-Poza and Sousa-Poza, 2000; Wanous, Reichers, and Hudy, 1997).² Job satisfaction is a reflective assessment of one's overall working environment that also incorporates expectations, norms, values, alternatives, and the outcomes and rewards of work (Angrave and Charlwood, 2015; Weiss,

²In the psychology literature, a common measure of subjective well-being is the multidimensional scale by Green, Felstead, Gallie, and Inanc (2016); Warr (2007, 1990), which comprises enthusiasm/depression, and anxiety/comfort dimensions. Meanwhile, single-item measures of job satisfaction are typically as valid and reliable as their multi-item counterparts (Wanous et al., 1997).

2002). The fact that job satisfaction predicts labour market behaviours such as job quits (Clark, 2001; Green, 2010; Lévy-Garboua, Montmarquette, and Simonnet, 2007) implies that this measure reflects workers' preferences (Clark, 2015). Job satisfaction is also instrumentally important for productivity (Böckerman and Ilmakunnas, 2012). The job satisfaction literature has examined how different working conditions and arrangements influence job satisfaction. For example, analyses of the 1997 International Social Survey Program show that an interesting job and good relations with management are the biggest predictors of job satisfaction (Clark, 2005; Sousa-Poza and Sousa-Poza, 2000). Similarly, relying on German panel data, Cornelißen (2009) identifies relationships with colleagues and managers, job insecurity, and task diversity as the most influential satisfaction determinants. Clark (2011) proposes that job satisfaction is a comprehensive summative measure reflecting objective and subjective job quality, though this view has not remained unchallenged (Brown et al., 2012). The main critique of using job satisfaction as a proxy for worker well-being is that workers may report being satisfied with jobs that are objectively bad. Specifically, individuals may adapt to low job quality and learn to be satisfied with poor working conditions (Brown et al., 2012). Therefore, job satisfaction reflects both well-being at work and the norms and expectations that employees have when answering such questions, which is something that should be kept in mind when interpreting these data (Brown et al., 2012). Nevertheless, subjective measures provide a valuable bottom-up perspective on workers' own understanding of well-being. The objective and subjective approaches are not mutually exclusive and are often used in complementary ways (Green, 2006; Green et al., 2016).

In this paper, we study self-reported perceptions of work meaningfulness and as such, we draw on the subjective approach, while also recognising the critiques and insights of its objective counterpart. A key limitation of both the objective and subjective approaches to job quality is the lack of attention to work as a source of meaning.³ This general neglect of work meaningfulness is unfortunate, because it severely limits our understanding of the true spectrum of work well-being measures and the particular position of meaningful work in that spectrum. It is also surprising, given that the notion of work meaningfulness is not new in the social sciences and can be traced back at least to Karl Marx who believed that labour was inherently purposeful and a source of fulfillment, rather than just a means to satisfy material needs (Spencer, 2009).⁴

³To our knowledge, the only notable exception is the newly-created job quality index by the Gallup organisation, which includes having a sense of purpose and dignity at work as one of the ten job quality indicators (Rothwell and Crabtree, 2019).

⁴According to Marx, capitalism eroded people's ability for self-actualisation and control over their work –

Nevertheless, the fact that self-reported measures of job satisfaction are well-established raises the question of the value-added of research on meaningful work perceptions. Indeed, job satisfaction captures the overall subjective evaluation of the working environment. However, the concept of work meaningfulness goes above and beyond job satisfaction. For example, a person can be dissatisfied with the general working conditions, and find their daily duties stressful and unpleasant, yet deem the nature of the tasks as meaningful or impactful. Individuals working in occupations involving teaching or nursing easily fit this description. Conversely, an individual can be satisfied with the working conditions on the job but still perceive their work activities as meaningless. This may explain why, for example, many people do not quit their jobs despite finding them socially useless (Dur and van Lent, 2019). Empirically, Allan, Batz-Barbarich, Sterling, and Tay (2019) provide evidence that different scales of meaningful work are correlated with but distinct from job satisfaction. We find similar results with our data: the correlation between meaningful work and satisfaction with working conditions is 0.33 (see Table 4.A.5).

Therefore, like Steger, Dik, and Duffy (2012), we argue that meaningful work is a *eudaimonic* dimension of (perceived) worker well-being. Eudaimonia generally entails flourishing and living a life that realises one's potential (Deci and Ryan, 2008; Graham and Nikolova, 2015; Ryan, Huta, and Deci, 2008), and contrasts with *hedonic* and *evaluative* dimensions of subjective well-being. Typically, eudaimonic subjective well-being is captured using survey questions about whether the respondent has meaning and purpose in life. For example, Graham and Nikolova (2015) find that the biggest predictor of eudaimonic well-being is belief in hard work as a means of getting ahead in life, highlighting the connection between efforts in the work domain and having a life purpose. By contrast, life satisfaction is an evaluative measure of well-being and is a reflective assessment of one's overall life circumstances. Therefore, just like life satisfaction is distinct from having meaning and purpose in life (Graham and Nikolova, 2015), job satisfaction is conceptually different from work meaningfulness. Evaluative and eudaimonic measures also differ from hedonic well-being, i.e. positive and negative feelings at a particular point in time, such as stress or anger, or happiness and joy triggered by specific events (Graham et al., 2018). Hedonic well-being in the workplace refers to feelings of stress, engagement, and enthusiasm. Table 4.A.5 demonstrates that the correlations between meaningful work and stress, engagement, and enthusiasm range between 0.1 and 0.4, highlighting the difference between hedonic and eudaimonic subjective well-being at work.

a process known as "alienation" (Spencer, 2009, 2015).

4.2.2 The preconditions for meaningful work

Our understanding of meaningful work is based on self-determination theory (Deci and Ryan, 1985; Ryan and Deci, 2000b). According to this theory, satisfying three innate psychological needs —competence, autonomy and relatedness —underpins intrinsic motivation and eudaimonic well-being (Ryan and Deci, 2000b). Without competence, autonomy, and relatedness, individuals are unable to derive utility from meaning. This justifies the conceptual and empirical examination of work meaningfulness in the context of these three preconditions.

In self-determination theory, *competence* refers to the perceived ability to successfully overcome challenging tasks at work and contribute to a cause, which creates feelings of mastery (Martela and Riekk, 2018; Rosso et al., 2010; Ryan and Deci, 2000b). It entails a belief in having the right skills to make an impact. Moreover, people satisfy their need for *autonomy* when they perceive that they have choices and authority over what to do. Autonomy is empirically linked to meaningfulness (Martela and Riekk, 2018; Martela, Ryan, and Steger, 2018; Ryan and Deci, 2000b) because it allows for self-expression, control over the work content and process, and the ability to choose how and when to apply different skills and capabilities. This is in stark contrast with *heteronomy*, a condition whereby behaviour is regulated by forces that the worker perceives as over-imposed, as would be the case with heavy top-down management, for example. Finally, *relatedness* is about the inter-personal relationships at work (Martela and Riekk, 2018; Ryan and Deci, 2000b). Workers feel related if they experience genuine care from their bosses or colleagues, and that they care about their superiors and co-workers in return.

Importantly, autonomy, competence, and relatedness are not externally determined objective targets, but rather strongly depend on each individual's innate needs. This implies that there is no single policy in the workplace that employers can adopt to meet the needs of all employees. In addition to autonomy, competence, and relatedness, environmental circumstances and extrinsic rewards facilitate or forestall self-motivation. Therefore, motivation is formed through the interplay between the work environment created by the employer and the satisfaction of the person-specific needs for autonomy, relatedness and competence.

According to Ryan and Deci (2000b), there is an intricate relationship between different states of motivation and meaning. First, when work completely fails to satisfy people's innate needs, they are *amotivated*, meaning that they are passive and unwilling to work at any level of pay. Second, when workers are in a state of *controlled motivation*, their needs for autonomy, relatedness, and competence are

partially satisfied. However, in this state, workers can only be extrinsically motivated through monetary rewards because they do not find their tasks inherently purposeful. Finally, workers are in a state of *autonomous motivation* when their psychological needs are fulfilled and they feel that the purpose of their tasks matches their personal values and purpose. At that point, their tasks become meaningful. Autonomous motivation is impossible at low levels of autonomy, competence, and relatedness, because people cannot experience self-efficacy: they fail to see how their personal actions affect the outcome, which implies that their effort is meaningless. Therefore, the three basic needs for autonomy, competence, and relatedness, need to be satisfied to derive meaning from work.

Extrinsic factors such as financial incentives and rewards may be additional preconditions to achieving work meaningfulness (Cassar and Meier, 2018; Spencer, 2015). For instance, an insufficient income level limits the ability to meet basic consumption needs and thwarts people's efforts to achieve their goals and work independently (Vohs, Mead, and Goode, 2006). Financial incentives matter for intrinsic motivation and effort (Lazear, 2018), and psychologists show that pecuniary aspects are more important than non-pecuniary ones to workers who have less income to begin with (Rosso et al., 2010). The flipside of this argument is that jobs deprived of intrinsic rewards would only matter to workers through the monetary compensation they offer (Cassar and Meier, 2018). Other extrinsic factors—hours of work, career progression possibilities, and job insecurity—also influence the ability to derive work meaningfulness (Spencer, 2015). For instance, individuals will work longer or more intensely if they feel motivated, compared to a state when they are only working to finance their consumption and leisure (Cassar and Meier, 2018). Nevertheless, long working hours can lead to exhaustion and limit the ability for creative work. Finally, job insecurity can negatively affect certain health outcomes (Caroli and Godard, 2016), including mental health (Reichert and Tauchmann, 2017), which may leave little scope for meaningfulness.

4.2.3 Meaning and labour economics: an overview of the literature

The extant literature provides some intuition into how meaningful work perceptions can affect workers' choices and behaviours. First, experimental studies have documented that viewing one's work as meaningful (i.e. task meaning) increases effort and productivity (Ariely, Kamenica, and Prelec, 2008; Bäker and Mechtel, 2018; Chadi, Jeworrek, and Mertins, 2016; Chandler and Kapelner, 2013; Grant, 2008; Kosfeld et al., 2017). For example, Ariely et al. (2008) conducted two experiments manipulating the meaningfulness of the task (finding consecutive occurrences of the letter 's' in Experi-

ment 1 and assembling Legos in Experiment 2) by acknowledging or destroying the final product, which revealed to participants its meaningfulness. Indeed, labour supply was higher and reservation wages were lower when the experimenters signalled the meaningfulness of the task. Second, theoretical work in economics has proposed that people have an innate drive for sense-making (Chater and Loewenstein, 2016; Karlsson, Loewenstein, and McCafferty, 2004). Thus, fulfilling the need for meaning is part of an individual's utility function and decision-making.

Our paper is related to but substantively different from four recent contributions. First, using pooled cross-sectional data from the International Social Survey Program (ISSP), Dur and van Lent (2019) document that about one in ten employees finds their job useless, with the share being the highest among those engaging in routine tasks as well as those in sales, finance, public relations, and marketing. The authors furnish suggestive evidence that the share of socially useless jobs may increase with the output gap, which they interpret as evidence for labour hoarding, i.e. the retention of more workers than necessary in times of economic turmoil. Dur and van Lent's (2019) study differs from ours in that it focuses on the determinants of a different concept—socially useless jobs—and utilizes a different dataset.

Several studies have relied on time use data to study meaningfulness during work episodes. For example, Kaplan and Schulhofer-Wohl (2018) rely on the American Time Use Survey and find that since the 1950s, women have switched to occupations that have provided high non-pecuniary benefits of working, including more episodes of meaningfulness. The conclusion is more nuanced for men who have shifted to occupations that generate less meaningfulness and happiness and more stress, but fewer feelings of pain and tiredness. Nevertheless, Bryce (2018) surprisingly documents that while working is in itself negatively associated negatively associated with meaning, community and social service, legal, education, and healthcare occupations are considered the most meaningful (relative to transportation jobs) in the American Time Use Survey. The unexpected negative association between working and meaningfulness is likely due to the fact that time use surveys capture the *hedonic* and not *eudaimonic* work aspects.

Using German time use data, Wolf, Metzger, and Lucas (2019) document that along with childcare and exercising, working is among the most meaningful activities in people's daily lives, which contrasts with the findings of Bryce (2018). Wolf et al. (2019) find that individuals reporting no meaning at all and those reporting very high meaning are more likely to derive pleasure at work. They explain this seemingly paradoxical finding by noting that some individuals do not experience utility from meaning.

Our research fundamentally differs from these contributions. Specifically, we explicitly focus on meaningful work rather than related concepts such as socially useless work as in Dur and van Lent (2019). Moreover, we demonstrate that meaningful work determines relevant economic outcomes, such as retirement, training, and absenteeism. This implies that subjective evaluations of meaningful work relate to important behavioural consequences and are of interest to labour economists.

4.3 Data

We rely on three waves of the European Working Conditions Survey (EWCS), conducted in 2005, 2010, and 2015 (European Foundation for the Improvement of Living and Working Conditions, 2019, 2007, 2012, 2017). The EWCS is a well-known data source for studying the well-being implications of working conditions (see, for example, Aleksynska (2018), Caroli and Godard (2016), Cottini and Lucifora (2013)). While the survey is performed every five years since 1990/1, our analysis is constrained to the 2005-2015 waves due to the availability of key variables for our research. We note that the related dataset—International Social Survey Programme (ISSP)—which included work orientations modules in 1989, 1997, 2005, and 2015, asks respondents whether their job is useful to society (Dur and van Lent, 2019), but not about other aspects of engaging in meaningful work such as job fulfillment. Another disadvantage of this dataset is that the country coverage has changed over time.

Our final analysis sample focuses on the common set of countries included in all three EWCS waves: the 28 EU Member States, Turkey, and Norway. We limit the analysis to this country set to ensure that our results are not driven by changes in the sample composition across the waves. In addition, this sample restriction is useful when implementing pseudo panel techniques.

For each wave, the survey polled about 1,000 individuals in each country. In some years, certain countries are over-sampled and therefore have a larger number of observations. While we have no control over the way in which the survey is conducted, in Model (1) of Table 4.3 we show that the differences in the number of observations across countries do not drive our main results.

The analysis sample comprises individuals who formally work part- or full-time. We exclude the unemployed and those out of the labour force. While our main analysis sample excludes respondents with missing information on any of the variables, in Appendix Table 4.A.4 we also provide analyses addressing any concerns related to bias arising from item non-response. The main analyses automatically exclude the self-employed as this group were not asked questions relating to relationships with

colleagues and superiors, permanent contracts, or benefits. Nevertheless, because the relationship between self-employment and meaningfulness is interesting, we offer additional analyses with the sub-set of available variables in Table 4.2. Table 4.A.1 in the Appendix details the number of observations per country and year in the final analysis sample (N=48,420).

4.4 Variables

4.4.1 Measuring meaningfulness

Shedding light on the causes and consequences of meaningful work requires appropriately measuring the underlying concept. At the outset, we acknowledge that this is a challenging task because of the lack of consensus on the concept's measurement. The most well-established and widely-used scale in the psychology literature is the Work and Meaning Inventory (WAMI) (Bailey et al., 2019a; Steger et al., 2012). WAMI captures aspects of positive meaning (i.e. having a career that one considers meaningful), meaning-making through work (i.e. work that helps the respondent to make sense of the world around him/her), and greater good motivations (i.e. having a job that is useful to society). Unfortunately, no nationally representative survey to date has included the WAMI questionnaire and it is unclear whether this scale is indeed valid and reliable in such contexts.

Ours is the first attempt to systematically study meaningful work perceptions using existing nationally-representative data and should be seen as the starting point for future investigations. Our meaningfulness of work measure is therefore based on available questions from the EWCS that most closely match the conceptual definition of meaningful work based on the literature.

Specifically, we identified one question reflecting greater good motivations – having the feeling of doing useful work. Second, we include a variable based on whether work is a source of fulfillment, i.e. whether the respondent's job gives the feeling of work well done. These two variables were also the basis of a summative meaningful work index used by Eurofound (2012) based on the 2010 EWCS. While the 2015 survey contains a question about not doubting the importance of one's work, which would have also been relevant to us, we choose these two variables as they are consistently available in all three survey waves that we use.

Figure 4.1 is a violin plot detailing the distribution of responses at each answer category for the two variables underlying the meaningful work index. A violin plot is a combination of box and density plots. The white dot represents the median and

the box denotes the interquartile range. The majority of responses and the medians for both variables are concentrated in the “strongly agree” or “agree” categories, which indicates low variation in the distribution of responses. About 5 percent of respondents disagree or strongly disagree about the usefulness of their job and work giving them the feeling of work well done.

We combine the two variables into an index by extracting the first component of a polychoric principal component analysis (PCA), which is a well-established data reduction procedure. Standard PCA assumes that the underlying data are continuous and normally distributed. As this is not the case for our variables, we relied on the polychoric version of PCA, which assumes that variables are ordered measurements of an underlying continuum and is therefore better suited for categorical variables (Olsson, 1979b). Polychoric PCA exploits the linear combinations of the polychoric correlation matrix of the input variables and preserves the ordinal or binary nature of the variables (Olsson, 1979b). The two variables that we use to create the index – the perceptions of doing useful work and a job well done — are relatively highly correlated: the simple Pearson’s correlation coefficient is 0.6 and Cronbach’s $\alpha=0.75$, which is a good starting point for using PCA. The first principal component accounts for 85 percent of the total variance (eigenvalue = 1.71). Meanwhile, the second component explained only 15 percent of the total variance and had an eigenvalue of 0.29. Following the Kaiser rule, we keep only the first principal component, which we rescale to range between 0 and 100 for ease of interpretation. The violin plot in Figure 4.2 demonstrates that the majority of observations are concentrated in the range between 75 to 100 and the median respondent in our sample is at about 90.

4.4.2 Independent variables

Following the conceptual framework outlined before, our key independent variables include measures of autonomy, competence, and relatedness, which we construct using polychoric PCA (see Table 4.A.2 in the Appendix).

Importantly, the measures of autonomy, relatedness, and competence do not reflect personal needs or the objective working conditions that employers have created, but rather reflect the interplay between the needs and the environment. Therefore, the self-reported index of autonomy reflects the match between the working conditions and the worker’s personal need for autonomy. As such, the measures for autonomy, competence, and relatedness reflect the degree to which the worker perceives their innate needs to be satisfied.

The autonomy index is based on variables capturing process and decision auton-

Figure 4.1: Violin plots: feelings of work well done and doing useful work

Source: Authors based on European Working Conditions Surveys (EWCS) 2005-2015

Figure 4.2: Violin plot, meaningful work index

Source: Authors based on European Working Conditions Surveys (EWCS) 2005-2015

omy, such as the freedom to take a break at will, change or choose tasks, methods of work, and the speed of work, conducting self-assessments, and applying one's own ideas at work. The Cronbach's α , the scale reliability coefficient is 0.66. The first principal component accounts for 56 percent of the total variance (eigenvalue = 3.36). Meanwhile, the second component explained only 15 percent of the total variance and had an eigenvalue of 0.91. Other components explained even less of the variance, which is why we only kept the first extracted principal component.

Our competence measure is based on workers' assessments of their skills, problem-solving, and learning. While the Cronbach's α of 0.4 is moderate, these variables capture the learning and self-efficacy aspects of self-determination theory. We extracted the first principal component after applying polychoric PCA (eigenvalue = 1.70), which explained 56 percent of the total variation. Finally, we construct a relatedness index based on variables indicating whether the respondent receives help and support from colleagues and their boss. While the wording of these questions is slightly different in the 2005 questionnaire, these variables were the only ones in the EWCS questionnaire reflecting relatedness ($\alpha = 0.73$). We extracted the first polychoric principal component (eigenvalue= 1.69), which accounted for 85 percent of the total variation.

In addition, we incorporate variables capturing monetary compensation and effort, as well as proxies for career advancement prospects and job insecurity, as discussed in Spencer (2015). Our analyses also feature standard socio-economic, demographic, and job controls such as age, gender, education, tenure, household size, presence of children in the household, marital status, whether the respondent is a public employee, the number of people the respondent supervises, company size, an indicator for having a permanent contract, and industry and occupation dummies. We do not include an indicator for part-time work due to the large number of missing observations in 2010. In addition, to account for meaningfulness spillovers, we control for whether the respondent has other jobs and whether they volunteer in their free time. We also include the following interview controls: interviewer fixed effects, the duration of the interview (in minutes), the presence of other people during the interview, and the interview month and day. All of these factors could influence subjective well-being responses (Conti and Pudney, 2011; Rehdanz and Maddison, 2005), which necessitates their inclusion in the regressions. Finally, we incorporate survey wave and country dummies, which adjust for any common temporal shocks (e.g. recessions) over time and differences in labour market features and institutions between countries, such as the presence of minimum wage laws and employment protection programmes. We adjust for cost-of-living differences in income across

countries by dividing income by the Eurostat's Purchasing Power Parity index (see Table 4.A.2). Appendix Table 4.A.3 provides summary statistics for key variables used in the analysis.

4.5 Empirical Approach

We model the perceived meaningfulness of work M of individual i living in country c in survey wave t as:

$$M_{ict} = \beta_0 + I_{ict}\gamma + E_{ict}\delta + Z_{ict}\alpha + S_{ict}\omega + \pi_c + \mu_t + \varepsilon_{ict} \quad (4.1)$$

where I is a vector of the preconditions for motivation based on self-determination theory (autonomy, competence, and relatedness) and E is a vector of extrinsic factors (income, benefits, working hours, job insecurity, and career prospects). Finally, Z is a vector of socio-demographic and job characteristics, such as gender, age, education, marital status, firm size, permanent contract, and others; π and μ denote country and year dummies, respectively; S is a vector of interviewer fixed effects and interview controls, and ε is the stochastic error term. Since the dependent variable ranges from 0 to 100, we estimate equation (4.1) using an Ordinary Least Squares (OLS) estimator.

In addition, given that we have repeated cross sections, in robustness checks we also implement a pseudo panel approach (Deaton, 1985), whereby we treat as a cohort workers sharing the same characteristics, such as birth year, age, or gender. The group-averages of these cohort variables are the new unit of analysis in the pseudo panel. The repeated cohort-level information implies that fixed effects estimators are possible. In this paper, we define the cohort based on age, gender, marital status, education, and country of residence. The results are reported in Table 4.3, Model (2).

Our results are correlational, as opposed to causal, for several reasons. First, individuals who value meaningful work likely self-select into jobs that provide intrinsic rewards and meaning. Therefore, traits such as intelligence, motivation, or pro-sociality could influence both job choice and meaningfulness perceptions. The pseudo panel strategy is an attempt to mitigate such concerns. Second, while intrinsic and extrinsic work rewards may influence meaningfulness, perceiving one's job as meaningful may influence effort, which in turn influences pay and intrinsic rewards. Ideally, we would have preferred to have individual-level panel data and exogenous variation in working conditions that would have allowed us to control for time-invariant unobserved heterogeneity and certain self-selection issues. Nevertheless, even though we only have pooled cross-sections, to the extent possible, we mitigate

endogeneity issues by including a large set of covariates, country dummies, and interview controls.

4.6 Results

4.6.1 Main results

Table 4.1 details our main results. Model (1), which is the basis for the computationally-intensive Shapley R^2 decomposition shown in Figure 4.4, includes a parsimonious set of controls. Model (2) is our baseline specification that features all key independent variables based on our conceptual framework, as well as the full set of socio-demographic controls, year and country fixed effects, and interview controls. Models (3) to (5) add interaction terms, which absorb additional variation in work meaningfulness and account for further heterogeneity. Model (3) considers the possibility of differences in meaningfulness perceptions across workers in the same industry, but in different occupations. In Model (4), the *education* $\times$ *occupation* dummies account for differences in meaningful work perceptions across people with the same level of education, but working in different occupations. In Model (5), we allow for the possibility of meaningfulness differences across people working in the same occupation, but living in different countries.

Table 4.1 demonstrates that both the magnitudes and the statistical significance of the coefficient estimates remain stable across these different specifications. Autonomy, relatedness, and competence are all positively associated with meaningfulness, whereby, for example, a 10-point increase in autonomy corresponds to an increase of 1.3-1.4 points in meaningfulness, which appears rather modest in magnitude. Nevertheless, autonomy accounts for a significant share of the variation in meaningfulness, as shown in Figure 4.4 below. Meanwhile, income and benefits are not associated with meaningfulness, which is an interesting result. The raw correlation coefficient between income and meaningfulness is also rather low ($\rho = 0.04$), which suggests that meaningfulness mostly reflects non-pecuniary work aspects. Future research should explore in greater detail the relationship between work meaningfulness and income and whether and to what extent income is a necessary precondition for motivation and meaning.

Perceptions of career advancement and job insecurity matter for meaningfulness in the expected directions, and longer working hours decrease meaningfulness, suggesting that excessive work intensity may limit the ability to derive work meaningfulness. Tenure, the number of working days, being a public employee, having

a permanent contract, working multiple jobs, and supervising others do not influence work meaningfulness, but respondents working in smaller firms have higher meaningfulness perceptions, compared to those working in larger firms. This finding may at first appear at odds with the positive relationship between relatedness and meaningfulness. Nevertheless, the negative coefficient on the firm size dummies is likely capturing aversion to hierarchy and preferences for autonomy (Benz & Frey, 2008).

Table 4.1: Determinants of meaningful work perceptions

	(1) No interview controls	(2) Interview controls	(3) Industry × occupation	(4) Education × occupation	(5) Country × occupation
Autonomy	0.138*** (0.004)	0.126*** (0.004)	0.127*** (0.004)	0.126*** (0.004)	0.127*** (0.004)
Competence	0.043*** (0.004)	0.039*** (0.004)	0.037*** (0.004)	0.039*** (0.004)	0.035*** (0.004)
Relatedness	0.192*** (0.004)	0.166*** (0.005)	0.166*** (0.004)	0.166*** (0.004)	0.165*** (0.004)
Log monthly income (PPP-adjusted)	0.059 (0.115)	0.269 (0.231)	0.174 (0.223)	0.290 (0.223)	0.126 (0.224)
Benefits & performance pay	-0.059 (0.189)	-0.007 (0.212)	-0.013 (0.210)	-0.004 (0.210)	-0.023 (0.210)
Job insecurity	-3.486*** (0.250)	-3.454*** (0.264)	-3.450*** (0.233)	-3.451*** (0.233)	-3.435*** (0.233)
Career advancement	3.902*** (0.171)	4.740*** (0.186)	4.700*** (0.191)	4.742*** (0.191)	4.648*** (0.191)
Log weekly hours	-2.574*** (0.269)	-2.332*** (0.321)	-2.387*** (0.312)	-2.343*** (0.312)	-1.882*** (0.315)
Weekly workdays (no.)		0.142 (0.134)	0.186 (0.127)	0.142 (0.127)	0.189 (0.128)
Public employee		0.616*** (0.235)	0.537** (0.239)	0.616*** (0.238)	0.527** (0.238)
<i>Firm size: Ref: 1-9 emp</i>					
10-249 employees		-1.072*** (0.207)	-1.122*** (0.203)	-1.068*** (0.203)	-1.082*** (0.203)
250 and more employees		-2.236*** (0.305)	-2.264*** (0.293)	-2.246*** (0.293)	-2.088*** (0.294)
Permanent contract		0.060 (0.262)	0.038 (0.247)	0.072 (0.247)	-0.057 (0.247)
Supervision (log no.)		0.157 (0.106)	0.176 (0.116)	0.156 (0.116)	0.151 (0.117)
Other jobs		-0.313 (0.307)	-0.320 (0.304)	-0.308 (0.305)	-0.247 (0.304)
Volunteer		-0.080	-0.070	-0.078	-0.079

Continued on next page

Table 4.1 – Continued from previous page

	(1)	(2)	(3)	(4)	(5)
		(0.196)	(0.196)	(0.196)	(0.196)
Tenure		0.020*	0.020*	0.021*	0.017
		(0.011)	(0.011)	(0.011)	(0.011)
Age		0.158***	0.158***	0.158***	0.163***
		(0.010)	(0.010)	(0.010)	(0.010)
Male		-0.296	-0.408**	-0.320*	-0.322*
		(0.200)	(0.195)	(0.194)	(0.194)
<i>Education. Ref: Low</i>					
Secondary education		-1.242***	-1.250***	-0.897	-1.342***
		(0.312)	(0.287)	(1.508)	(0.291)
Tertiary education		-2.821***	-2.758***	-4.073**	-2.826***
		(0.448)	(0.432)	(1.614)	(0.437)
Household size		0.102	0.106	0.103	0.131
		(0.093)	(0.088)	(0.088)	(0.088)
Spouse in household		0.569**	0.551***	0.572***	0.588***
		(0.222)	(0.212)	(0.212)	(0.212)
Children in household		0.249	0.223	0.252	0.333
		(0.212)	(0.206)	(0.206)	(0.205)
N	48,420	48,420	48,420	48,420	48,420
Adj.R ²	0.205	0.354	0.355	0.354	0.361

Source: Authors based on the European Working Conditions Surveys (EWCS) 2005-2015.

Notes: Robust standard errors in parentheses. The dependent variable in all models is perceptions of being engaged in meaningful work, which is an index measured on a scale of 0 to 100. See Table 4.A.2 for variable definitions. All regressions include occupation, industry, country, and year fixed effects. Models (2)-(5) also include interview controls (duration, number of people present during interview, interview month, and interview day, interviewer fixed effects), and individual control. Model (3) includes industry $\times$ occupation fixed effects, Model (4) includes education $\times$ occupation fixed effects, and Model (5) includes country $\times$ occupation fixed effects. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Furthermore, more educated respondents experience their jobs as *less* meaningful compared to workers with an elementary education, which is a finding worthy of further explorations. This seemingly paradoxical result is consistent with models of job crafting, according to which low-skilled individuals are able to see beyond their immediate tasks and find meaningfulness and purpose in seemingly menial tasks (Both-Nwabuwe, Dijkstra, and Beersma, 2017; Wrzesniewski and Dutton, 2001). Finally, women experience their jobs as more meaningful, compared to men with the same working conditions.

Model (2) of Table 4.1 also includes occupation and industry fixed effects. Figure 4.3 graphically summarises the differences in meaningfulness between occupations and industries, based on the regression coefficients from Model (2) in Table 4.1. Plant

Figure 4.3: Regression-adjusted differences in meaningfulness, by occupation and industry, with 95% confidence intervals

Source: Authors based on European Working Conditions Surveys (EWCS) 2005-2015

Notes: The figure shows the regression coefficients on the occupation and industry fixed effects from Table 4.1, Model (2). The reference category Panel A is managers and in Panel B: agriculture, hunting, forestry, and fisheries.

and machine operators, professionals, service and sales workers, and technicians find their jobs more meaningful compared to managers. The craft and related trades industry is the occupation with the greatest meaningfulness score, likely due to the creative nature of these jobs. In addition, workers in all industries except certain services industries find their jobs less meaningful than those working in the agricultural sector. In summary, the occupational and industry fixed effects point to a pattern whereby workers in the service industry and those in occupations providing creativity and autonomy tend to have greater meaningfulness perceptions.

Next, using Shapley-based decompositions (Israeli, 2007; Shorrocks, 2013) we empirically demonstrate the relative importance of key job characteristics for meaningful work. The Shapley-based decomposition method extracts the separate contribution to the explained variation in meaningfulness of each included independent variable. Specifically, Figure 4.4 indicates the relative contribution to the overall R^2 explained

Figure 4.4: Shapley-based decompositions

Source: Authors based on European Working Conditions Surveys (EWCS) 2005-2015

Notes: The figure shows the Shapley decompositions based on Model (1) in Table 4.1. $R^2 = 0.205$

by the different factors in Model (1) in Table 4.1. Our proxies for autonomy, competence, and relatedness account for 60 percent of the variation in meaningfulness of work. Income and benefits together account for less than half a percent. All in all, income, benefits, job insecurity, career advancement, and working hours explain about 13 percent of the variation in meaningfulness. The key insight from Figure 4.4 is that intrinsic rewards from work are about 4.6 times more important for meaningfulness compared with compensation and other extrinsic factors. Meanwhile, relatedness is the most important determinant of work meaningfulness.

Our main analyses exclude the self-employed because of a lack of information on the questions comprising the relatedness index, as well as those pertaining to benefits and performance pay and permanent contracts. Nevertheless, the self-employed enjoy greater mental health and subjective well-being compared to similar regular employees (Benz and Frey, 2008; Binder and Coad, 2013; Blanchflower and Oswald, 1998; Hessels, Arampatzis, van der Zwan, and Burger, 2018; Nikolova and Graham, 2014; Nikolova, 2019). This well-being premium is often attributed to the utility of being your own boss and having autonomy and flexibility (Benz and Frey, 2008;

Table 4.2: Meaningful work and self-employment

	(1)	(2)	(3)
		Sub-sample	
	Self-employed control	Self- employed	Non Self- employed
Autonomy	0.149*** (0.003)	0.263*** (0.018)	0.144*** (0.003)
Competence	0.046*** (0.004)	0.038*** (0.015)	0.047*** (0.004)
Self-employed	1.520*** (0.294)		
Log monthly income (PPP-adjusted)	0.712*** (0.183)	1.411*** (0.482)	0.451** (0.216)
Job insecurity	-3.817*** (0.214)	-3.910*** (0.931)	-4.007*** (0.227)
Career advancement	5.542*** (0.176)	3.871*** (0.683)	5.750*** (0.189)
Log weekly hours	-2.148*** (0.262)	-1.202 (0.732)	-2.351*** (0.302)
N	57,867	6,661	51,206
Adj.R ²	0.323	0.347	0.322

Source: Authors based on the European Working Conditions Surveys (EWCS) 2005-2015.

Notes: Robust standard errors in parentheses. The dependent variable in all models is perceptions of being engaged in meaningful work, which is an index measured on a scale of 0 to 100. See Table 4.A.2 for variable definitions. All regressions include country and year fixed effects, interview controls (duration, number of people present during interview, interview month, and interview day, interviewer fixed effects), individual controls, and occupation and industry fixed effects. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Hyttinen, Ilmakunnas, and Toivanen, 2013). More recently, Wolfe and Patel (2019) demonstrate that, rather contradictorily, the self-employed are slightly more likely to perceive their jobs as socially useful, but are not more likely than regular employees to rate their work as important. These differences are likely due to the differences in sample composition in the two regressions used as the authors rely on both the ISSP and the EWCS data.

In light of these studies, we explore the relationship between self-employment and work meaningfulness by omitting the control variables that are not available for the self-employed sample, namely the relatedness index, benefits and performance pay, and permanent contracts. To our knowledge, this is the first exploration of work meaningfulness differences related to self-employment. Model (1) in Table 4.2 demonstrates that the self-employed enjoy higher levels of work meaningfulness, compared to private- and public-sector employees with similar working conditions and autonomy and competence levels.

Furthermore, in Models (2) and (3) we explore whether autonomy and competence matter more for the self-employed compared to private and public employees.

Both autonomy and income seem to have stronger associations with work meaningfulness in the self-employed sample, yet career advancement possibilities and working hours are more strongly associated with meaningfulness for the non-self-employed group. Income is statistically significant in these regressions likely due to the omission of the performance pay and benefits variable. It also matters more for the meaningfulness of the self-employed.

4.6.2 Robustness checks

Even though our estimation strategy only allows us to show conditional correlations rather than causal results, in Table 4.3 we present several robustness checks, which increase confidence in our results and conclusions. First, in Model (1), we adjust for the possibility that the results are driven by differences in the sample sizes across countries. Specifically, we re-estimate our main regressions using the inverse of the number of observations per country as a weight. The results remain virtually unaltered compared to those in Model (2) of Table 4.1, with the only notable difference being the marginally statistically significant coefficient estimate for income.

Furthermore, in Model (2) of Table 4.3, we create a pseudo panel whereby the level of analysis is a cohort comprised of respondents of the same age group, gender, marital status, education level, and living in the same country. The results are very similar to our baseline specifications. While we do not have a panel data set with observations on the same individuals followed over time, the pseudo panel findings provide suggestive evidence that our main conclusions will likely hold in a panel setting as well.

Next, in Model (3), we report the results using the simple average of the variables comprising the meaningful work index. The results are virtually identical to the main findings in Model (2) in Table 4.1, likely due to the fact that we only use two variables to construct the index. In the case of two inputs, the PCA and the simple average of the inputs often give similar results. In Model (4), we also address the concern that satisfaction with working conditions already captures objective and subjective working conditions, which may render work meaningfulness superfluous. If this was true, controlling for job satisfaction would yield the coefficient estimates on the key independent variables statistically insignificant. Model (4) in Table 4.3 demonstrates that this is not the case. Therefore, autonomy, relatedness, and competence matter for work meaningfulness above and beyond job satisfaction.

Finally, Models (5) and (6) of Table 4.3 differentiate between respondents who started a new job in the past two years vs. those working in the same firm for at

Table 4.3: Robustness checks

	(1) Weighted	(2) Pseudo Panel	(3) Alternative DV	(4) Working conditions satisfaction control	(5) At least 2 years on the job	(6) Less than 2 years on the job
Autonomy	0.127*** (0.004)	0.132*** (0.022)	0.126*** (0.004)	0.103*** (0.004)	0.121*** (0.004)	0.158*** (0.015)
Competence	0.042*** (0.004)	0.055** (0.023)	0.039*** (0.004)	0.044*** (0.004)	0.034*** (0.005)	0.047*** (0.016)
Relatedness	0.163*** (0.005)	0.208*** (0.021)	0.166*** (0.005)	0.127*** (0.004)	0.161*** (0.005)	0.184*** (0.018)
Log monthly income (PPP-adjusted)	0.420* (0.249)	0.128 (0.370)	0.269 (0.231)	-0.246 (0.226)	0.365 (0.258)	-0.615 (0.850)
Benefits & performance pay	-0.004 (0.228)	0.990 (1.291)	-0.007 (0.212)	-0.172 (0.206)	-0.015 (0.229)	0.597 (0.898)
Job insecurity	-3.464*** (0.282)	-6.200*** (1.345)	-3.456*** (0.264)	-1.957*** (0.258)	-3.416*** (0.303)	-3.308*** (0.862)
Career advancement	4.759*** (0.199)	6.305*** (1.204)	4.742*** (0.186)	3.196*** (0.183)	4.399*** (0.203)	5.982*** (0.783)
Log weekly hours	-2.293*** (0.336)	-1.501 (1.578)	-2.336*** (0.321)	-1.408*** (0.308)	-2.349*** (0.355)	-2.917** (1.163)
<i>Working conditions satisf.:</i>						
<i>Ref: not at all satisfied</i>						
Not very satisfied				7.357*** (0.774)		
Somewhat satisfied				14.238*** (0.750)		
Very satisfied				19.588*** (0.769)		
N	48,420	2,776	48,420	48,284	40,427	7,993
R ²	0.461	0.252	0.451	0.480	0.468	0.706

Source: Authors based on the European Working Conditions Surveys (EWCS) 2005-2015.

Notes: Robust standard errors in parentheses. The dependent variable in Models (1), (2), and (4), is perceptions of being engaged in meaningful work, which is an index measured on a scale of 0 to 100. See Table 4.A.2 for variable definitions. All regressions include country and year fixed effects, interview controls (duration, number of people present during interview, interview month, and interview day, interviewer fixed effects), individual controls, and occupation and industry fixed effects. Model (1) is a weighted regression using the inverse of the number of observations per country as a weight. Model (2) is based on a pseudo panel whereby the unit of observation is a cohort, defined as people in the same age group, gender, education level, marital status, and country. Model (3) uses a different dependent variable - the simple average of feeling of work well done and feeling of doing useful work, rescaled to range between 0 and 100. Model (4) controls for satisfaction with working conditions and Models (5) and (6) split the sample according to the respondent's duration of employment with the current employer. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

least 2 years. Through these specifications we test whether there is a honeymoon effect after a job switch (Chadi and Hetschko, 2018; Georgellis and Yusuf, 2016), whereby respondents perceive their jobs as meaningful due to the excitement related to the new job (i.e. the “magic of the new”) rather than the actual working conditions. Comparing the coefficient estimates in Models (5) and (6) indicates that the associations between autonomy, competence, and relatedness are slightly stronger for new employees, providing some support of the “magic of the new” hypothesis. Nevertheless, the differences between Models (5) and (6) and those in the main specification appear minimal.

In an additional robustness check shown in Appendix Table 4.A.4, we also address the issue of loss of information and potential bias arising from dropping observations with missing information. Specifically, we create an additional category for missing information for all control variables except autonomy, relatedness, and competence. Where the original variable is continuous, we create quartiles, and treat the variable as categorical with missing observations being the fifth category. The missing category for these variables has no particular interpretation but only serves to preserve the number of observations. The results demonstrate that the main patterns we identify in Model (2) of Table 4.1 still hold when we account for missing observations.

4.6.3 The labour market consequences of work meaningfulness

In this section, we demonstrate that meaningfulness perceptions are important for economists because they predict labour market behaviours in expected ways. Specifically, we estimate the relationship between perceptions of meaningful work and the number of sick days, the probability of reporting working when sick, the likelihood of participating in training, and the desired retirement age. As such, we provide the first validation of meaningful work perceptions and highlight their usefulness for labour economists. Ideally, in line with the job satisfaction literature, we would have tested how well meaningful work perceptions predict actual or intended job quits (Böckerman and Ilmakunnas, 2009; Clark, 2001; D’Ambrosio, Clark, and Barazzetta, 2018; Green, 2010; Lévy-Garboua et al., 2007). For example, Dur and van Lent (2019) show that individuals with socially useless jobs are more likely to report that they would like to change their jobs if they had the opportunity. Unfortunately, the EWCS lacks information on actual and intended job quits.

Table 4.4 details the results. First, in Model (1), we show that individuals who perceive their work as meaningful are likely to report fewer sick days. By exponentiating the coefficient estimate of 0.004, we find that a ten-point increase in the meaningfulness index corresponds to a decrease in the number of sick days by 10

Table 4.4: Meaningful work as a predictor of labour market outcomes

	(1) Log sick days	(2) Work when sick	(3) Participate in training	(4) Retirement age
Perceptions of meaningful work	-0.004*** (0.000)	-0.002*** (0.000)	0.001*** (0.000)	0.025*** (0.004)
Mean dependent variable	8.172	0.486	0.553	63.109
N	40,564	29,952	46,493	17,543
R ² /Pseudo R ²	0.338	0.162	0.257	0.422

Source: Authors based on the European Working Conditions Surveys (EWCS) 2005-2015.

Notes: Robust standard errors in parentheses. The dependent variable in Model (1) is the natural logarithm of the number of days the respondent was sick and absent from work in the past year; in Model (2), it is the probability of reporting to have worked while sick in the past year. This variable is only available for 2010 and 2015. In Model (3), it is the probability that the respondent participated in skills training in the previous year. Both model (2) and (3) report average marginal effects. In Model (4), the dependent variable is the age at which the respondent wishes to retire, whereby respondents who reported that they would like to keep working as late as possible are coded as wanting to retire at "80". Information for this variable is only available for 2015. The controls included are the same as Model (2) in Table 4.1. See Table 4.A.2 for variable definitions. Models (2) and (3) are estimated using a logistic regression and the average marginal effects are reported. Models (1) and (4) are estimated using OLS. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

percent. Evaluated at the mean of about 8 days, a ten-point increase in work meaningfulness reduces absenteeism by almost one day a year. Nevertheless, respondents who find their work more meaningful are in fact less likely to work when sick. A ten-point increase in meaningfulness decreases the probability of working while sick by 2 percent (from a baseline probability of about 50 percent). Taken together, these results suggest that meaningful work increases effort (through reducing absenteeism), but not at the cost of damaging health (i.e. workers still remain home when sick).

Furthermore, in Model (3), we demonstrate that a ten-point increase in meaningfulness corresponds to a 1 percent increase in the likelihood of participating in skills training programs. Finally, Model (4) shows that a ten-point increase in meaningfulness corresponds to a 2.5-year increase in the desired retirement age. The average intended retirement age in the sample is 63.1 years, which suggests that a ten-point increase in work meaningfulness could extend the age at which individuals wish to leave the workforce to about 66. While relatively small in magnitude, these results have important implications for policymakers faced with rising life expectancy and a greater share of older workers in the economy.

4.7 A Meaningful framework

The results presented above suggest that perceptions of meaningful work explain workers' labour market choices. Specifically, individuals who experience their jobs as more meaningful exert more effort, as evidenced by the lower absenteeism, higher

willingness to participate in training, and delayed retirement. Moreover, we find evidence that motivation arising from autonomy, relatedness, and competence, is relatively more important as a meaningfulness precondition compared to standard objective working conditions. In this section, we build on the conceptual insights of self-determination theory, the related literature on meaningful work, and our own empirical findings to derive a framework of motivation, meaning, and utility. Our framework can guide future explorations of meaningful work in labour economics.

4.7.1 A model of motivation and meaning

Standard neoclassical labour supply models view work effort as a disutility because it implies foregone leisure. Individuals will therefore work an additional hour if the compensation w is high enough to outweigh their disutility of exerting effort e . This “opportunity-cost view” assumes that employees only care about the monetary returns from work, and can only be motivated by extrinsic rewards, such as income and benefits. In contrast, self-determination theory posits that humans are motivated if they feel that their own actions directly impact their personal goals, i.e., when they experience self-efficacy (Bandura, 1997).

Building on Cassar and Meier (2018), we propose a simple theoretical framework that incorporates self-determination theory into utility theory. In our model, each worker’s utility function U depends on three arguments: utility from income Y , utility from meaning M , and utility from leisure, which negatively enters the model as the cost of exerting effort C . Income is a function of wages w and effort e , with decreasing marginal returns to effort. In this model, workers endogenously choose an effort level that equalizes the marginal benefits from income and the marginal costs of exerting effort.

We treat effort and motivation as two inputs in the utility function, where motivation is a vector of the satisfaction of the three psychological needs competence, relatedness and autonomy. We assume that motivation, which is formed through the interplay between the work environment created by the employer and person-specific needs for autonomy, relatedness and competence, is exogenous from the viewpoint of individual workers. For instance, an employer might grant equal levels of autonomy to all employees, but individuals will interpret these differently, depending on their own desire for autonomy. We assume that the psychological needs of workers are fixed over time, and that employers can create an environment in which their needs are satisfied. Employers’ increased investment in relatedness in the workplace by fostering stronger collegiality, for example, will therefore result in an exogenous increase in motivation.

To formalize, we introduce a cost-of-effort function that is related to both motivation and effort. In standard utility theory, the disutility of labour is always increasing in effort. We argue that when people are motivated, they have a lower disutility of effort. In other words, we argue that motivation reduces the cost of exerting effort in the sense that the reservation wage for exerting effort declines because the workers see that efforts have an impact on the outcome. Therefore, up to a certain point, motivated individuals experience a lower trade-off between work and leisure. As such, the cost function $C(e, a)$, is characterised by diminishing returns to motivation, such that $\frac{dC}{da} < 0$. Conversely, the cost of exerting effort is massive at low levels of motivation (amotivation), and will therefore have decreasing returns, such that $\frac{d^2C}{da^2} > 0$.

Furthermore, while the initial cost of effort is large, once a person has started exerting it, the marginal cost of extra effort diminishes ($\frac{dC}{de} < 0$ and $\frac{d^2C}{de^2} > 0$). In our model, both motivation and effort are constrained: a worker's ability to exert effort is not limitless. This is reflected by $e < E$, where E stands for the "burn-out level of effort". Similarly, motivation is capped at a level A , which is the level of complete autonomous motivation that requires no extrinsic rewards. We formalize this constraint as $a < A$.

Furthermore, we argue that work meaningfulness is only possible when the psychological needs are fulfilled and that self-efficacy increases meaningfulness (Ryan and Deci, 2000b). This is in line with our results presented in Table 4.1, where we show that the preconditions for motivation —autonomy, competence, and relatedness— are stronger predictors of meaningful work compared with objective working conditions. Therefore, we see work meaningfulness as a function of motivation a . The building blocks of our model give rise to the following utility function⁵:

$$U = Y(w, e) + M(a) - C(e, a) \quad (4.2)$$

where, as in Cassar and Meier (2018), $Y(w, e)$ is the utility from income as a means to consumption, which consists of monetary or extrinsic rewards w received for a certain level of effort e . In this setting, effort can reflect working hours or the intensity of work, for example.⁶ Income is positively related to exerted effort, but with diminishing returns such that $\frac{dY}{de} > 0$, $\frac{d^2Y}{de^2} < 0$. Workers will maximize utility by

⁵Note that a standard utility function would take the form $U(e) = Y(w, e) - C(e)$, where effort e increases income, but due to the theory of labour as disutility, decreases utility. As such, w needs to outweigh $C(e)$ in order for workers to supply labour

⁶Effort is conceptually different from productivity, as each person has a different level output for the same level of effort. Therefore, our simple model could be extended to include in Y a personal productivity parameter that determines the marginal output of extra effort.

endogenously choosing an effort level e , taking into account their level of motivation a . As a result, motivation will increase effort through the channel of decreased costs of exerting effort, which is important from an employer's viewpoint. If a increases, then $C(e, a)$ decreases, and thus the worker will increase his optimal level of effort e .

4.7.2 States of motivation

Next, we model the different states of motivation, namely, amotivation (passivity), controlled motivation, and autonomous motivation. Figure 4.5 graphically represents these states in light the functions for meaning M , income Y and cost C , plus total utility U . In the graph, the level of effort is fixed, to facilitate the discussion.

First, we argue that for each worker, there is threshold a^{min} reflecting the minimum level of motivation necessary to provide any effort and move out of passivity. This happens at the point where where $Y(w, e) = c(e, a)$, whereby the utility of income outweighs the costs of exerting effort. From this break-even point onward, people choose to work. Workers with motivation below this level a^{min} will not supply labour and will have a utility of 0.⁷

Second, we argue that once this level a^{min} is attained, people are *extrinsically* motivated by the wage they receive, as the utility of income outweighs the costs of exerting effort (because $M(a) \geq 0$). Graphically, this is represented in Figure 4.5 by the fact that the Y function lies above the U function. Third, we define a threshold level of motivation a^* whereby workers are able to derive meaning from their work. Reaching a^* , the point of autonomous motivation, requires the full satisfaction of the needs for autonomy, relatedness, and competence. Below this point, workers can exert effort as they are extrinsically motivated by their wages, but cannot derive utility from meaning.

In Figure 4.5, the transition between these states is graphically represented by the movement from amotivation ($a < a^{min}$), to passive compliance ($a^{min} < a < a^*$), to active personal commitment ($a \geq a^*$). The point a^* must differ from a^{min} as the innate needs for autonomy, competence, and relatedness must be at least minimally satisfied to trigger some form of action and move out of passivity. In addition, the threshold of a^* denotes the minimum level motivation needed to start deriving utility from meaning. These considerations give rise to the following piece-wise utility

⁷Note that in our dataset, we only observe workers, who by definition have $a > a^{min}$, which is a necessary condition for supplying labour.

Figure 4.5: Graphical Depiction of Conceptual Framework, Relating Motivation and Meaningfulness to Utility, for Fixed Levels of Effort

function describing the progression along the states of motivation:

$$U = \begin{cases} 0 & \text{if } a < a^{\min} \\ Y(w, e) - C(e, a) & \text{if } a^{\min} \leq a < a^* \\ Y(w, e) + M(a) - C(e, a) & \text{if } a \geq a^* \end{cases} \quad (4.3)$$

4.7.3 Behavioural consequences of motivation and meaning

We assume that workers sort into firms according to their expected motivation, i.e., the degree to which working conditions satisfy their needs for autonomy, competence, and relatedness. Based on this expectation, individuals decide how willing they are to work for a given firm, given a certain wage. In other words, this ‘reservation motivation,’ i.e., the minimum motivation needed to work for a certain wage is the driver of sorting into firms.

Furthermore, motivation is not fixed over time, is subject to changing policies at the firm level. Different policies can have different impacts on the experienced levels of motivation thus resulting in differing levels of effort throughout the employment duration, regardless of any firm or occupation switches. Potentially, changing motivation could trigger job quits or occupational switches, which is an empirical question that future research should investigate. Our model has two major behavioural consequences that can help guide further explorations of meaningful work in economics. First, it explains that some people may exert more effort than others for the same wage level due to differences in utility from meaning. In fact, this implication follows the logic of standard reservation wages and compensating differentials explanations. We also show this empirically by relating meaningfulness to measures of effort, such as absenteeism, participation in training, and preferences over delayed retirement. Therefore, work meaningfulness has large implications for worker behavior through the channel of motivation.

Second, our model can shed light on how the level of motivation gives rise to different preferences for meaning. Empirically, there is evidence for such heterogeneity in preferences (Fehrler and Kosfeld, 2014; Wolf et al., 2019). The model shows that this heterogeneity might be driven by differences in motivation, that causes some people to be unable to enjoy utility from meaning. Specifically, when motivation is below a^* , deriving utility from meaning is impossible. This may be empirically revealed by survey answers related to preference for meaning.

4.8 Conclusion and avenues for future research

We are the first to empirically investigate the factors influencing meaningful work using nationally representative samples of working-age adults in 30 European countries. Using data from the 2005, 2010, and 2015 European Working Conditions Surveys, our findings show that the non-monetary aspects of work, such as relatedness, autonomy, and competence, have a 4.6 times stronger association with the meaningfulness of work than income, job insecurity, benefits, and working hours. Importantly, we demonstrate that work meaningfulness predicts workers' effort, as measured by absenteeism, skills training, and retirement intentions. As such, we identify perceptions of having meaningful work as an important complement to extant measures of job quality.

Our study provides the first insights on the topic of work meaningfulness in economics. As such, it opens an exciting new research agenda, which, in our view, should prioritise three aspects. First, this paper's insights could inform the development of theoretical models that formally integrate self-determination theory into utility functions, which can guide future explorations of work meaningfulness in economics. In this chapter, we have proposed such a theoretical framework that extends the insights in Cassar and Meier (2018) by formally incorporating self-determination theory. Cassar and Meier (2018) follow the classical opportunity-cost view of labour and model individuals as deriving utility from both meaning and income and experiencing disutility from exerting effort. Each worker maximises utility by choosing an optimal level of effort (Cassar and Meier, 2018). In contrast, self-determination theory posits that individuals are motivated when they feel that their own actions directly impact their personal goals, i.e. when they experience self-efficacy (Ryan and Deci, 2000b). Therefore, a model of meaningful work should take into account that the disutility of exerting effort is decreasing in motivation: the cost of effort is lower for more motivated workers. Furthermore, according to self-determination theory, there is a minimum level of motivation required to experience any utility from meaning. This threshold is a crucial element in modelling meaningful work; its omission severely limits our understanding of how and why workers might make decisions based on the meaningfulness of their work.

Second, collecting and analysing longitudinal information on meaningful work perceptions is a logical extension of our research. A major advantage of the longitudinal design will be the repeated information on the same individuals over time, which should net out the influence of reporting bias in answering meaningful work questions as well as the influence of time-invariant norms and expectations. A panel

dataset will also be helpful in studying the short-term and long-term stability and consistency in responses to meaningful work questions within and across individuals. Moreover, it will facilitate the exploration of whether current meaningful work levels predict labour market behavior in subsequent survey waves.

Third, future research should investigate the interplay between the role of norms and expectations on the one hand and changing working conditions on the other for explaining variation in work meaningfulness. While self-reported meaningful work answers in part reflect these norms and expectations, unpacking the role of norms and expectations from that of actual working conditions is a crucial next step. Brown et al. (2012) provide two approaches to deepening our understanding of the role of norms and expectations in job satisfaction research, which can be applied to the study of work meaningfulness. First, they suggest complementing econometric analyses of job satisfaction with qualitative interviews about the role of extrinsic and intrinsic factors for job satisfaction answers. Second, the authors recommend explicitly controlling for norms and expectations in regression analyses by including variables measuring work orientations and job values. This is likely to be a viable way forward, yet it will be contingent on the clear conceptualisation and measurement of work orientations.

In short, we envision that our contribution will inspire a new line of research into the causes and consequences of meaningful work. This research agenda can provide timely novel insights into how to organise the future of work in a meaningful and dignifying way at a juncture that the future of work is in flux. Meaningful work is becoming increasingly salient in light of the ongoing processes of automation and digitalisation, which are altering the nature of paid and unpaid work activities. Against this backdrop, understanding what job characteristics enhance or diminish meaningfulness can provide important guidance to policy-makers and employers regarding boosting organisational performance and social functioning. Specifically, previous research shows that meaningful work is associated with higher productivity and lower turnover (Ariely et al., 2008; Rosso et al., 2010). In addition, this paper shows that those engaged in meaningful work are likely to remain longer in the workforce, which has implications for health and well-being and can help solve current demographic challenges related to ageing populations and rising dependency ratios (Nikolova and Graham, 2014). We also demonstrate that meaningful work can increase effort through reducing absenteeism and increasing the likelihood of participating in skills training.

By furnishing not only material means, but also social identity and individual self-esteem, work is a pivotal part of human life. Since most adults spend a large part

of their waking hours in work-related activities, understanding what factors make work a life-enriching and dignifying experience or, on the contrary, a degrading and meaningless one, can help design policies to enhance workers' well-being, boost organisational performance, and increase civic engagement and social welfare. Our findings underscore the importance of intrinsic factors for meaningfulness. Objective working conditions related to hierarchy, job insecurity, and working hours can create an important foundation enabling workers to gain meaningfulness from their jobs. However, it is autonomy, competence, and especially relationships at work that nourish and sustain meaningfulness. Future research should prioritise exploration of employer policies to encourage the satisfaction of these three innate needs, to promote meaningfulness in the workplace.

4.A Descriptive tables and figures

Table 4.A.1: Number of observations per country and year in the main analysis sample

Country	2005	2010	2015
Austria	307	401	465
Belgium	469	1,591	1,214
Bulgaria	539	461	496
Croatia	445	558	422
Cyprus	346	499	536
Czech Republic	347	431	414
Denmark	601	734	601
Estonia	287	486	457
Finland	577	487	522
France	387	1,238	874
Germany	384	991	775
Greece	448	396	300
Hungary	595	598	253
Ireland	502	430	472
Italy	387	518	347
Latvia	513	537	353
Lithuania	376	406	467
Luxembourg	290	324	453
Malta	342	499	606
Netherlands	559	523	488
Poland	470	584	395
Portugal	514	446	321
Romania	372	366	381
Slovakia	541	482	414
Slovenia	324	859	816
Spain	404	423	1,254
Sweden	682	495	586
UK	378	613	806
Turkey	321	882	799
Norway	558	664	672

Source: Authors based on the European Working Conditions Surveys (EWCS) 2005-2015

Table 4.A.2: Variable definitions

Variables	Explanation
<i>Dependent variable</i>	
Meaningful work index	A summary index created by extracting the first component of a polychoric principal component analysis (PCA) based on the statements (1) "your job gives you the feeling of work well-done" and (2) "you have the feeling of doing useful work." Both variables are originally measured on a scale, whereby 1=Never, 2=Rarely, 3=Sometimes, 4=Most of the time, 5=Always. The index is rescaled to range between 0-100 (higher score indicates a higher degree of perceived meaningfulness).
<i>Key independent variables</i>	
Autonomy	A summary index created by extracting the first component of a polychoric principal component analysis (PCA) of the following variables: (1) able to choose or change order of tasks, (2) able to choose or change methods of work, (3) able to choose or change speed or rate of work, (4) main paid job involves "assessing yourself the quality of your own work," (5) "you can take a break when you wish," (6) "you are able to apply your own ideas in your work." Variables (1)-(4) are originally measured on a scale 0=No, 1=Yes. Variables (5)-(6) are originally measured on a scale, whereby 1=Never, 2=Rarely, 3=Sometimes, 4=Most of the time, 5=Always. The index is rescaled to range between 0-100 (higher score indicates a higher degree of perceived autonomy).
Competence	A summary index created by extracting the first component of a polychoric principal component analysis (PCA) of the following variables: (1) respondent has appropriate skills to cope with current or more demanding duties, (2) main paid job involves "solving unforeseen problems on your own," (3) main paid job involves "learning new things." Variable (1) is originally measured on a scale 0=No, 1=Yes. Variables (2)-(3) are originally measured on a scale, whereby 1=Never, 2=Rarely, 3=Sometimes, 4=Most of the time, 5=Always. The index is rescaled to range between 0-100 (higher score indicates a higher degree of perceived competence).

Continued on next page

Table 4.A.2 – Continued from previous page

Variables	Explanation
Relatedness	<p>A summary index created by extracting the first component of a polychoric principal component analysis (PCA) of the following variables: (1) "your colleagues help and support you", (2) "your manager helps and supports you." Both variables are originally measured on a scale, whereby 1=Never, 2=Rarely, 3=Sometimes, 4=Most of the time, 5=Always. In 2005, the questions underlying these variables had a slightly different wording, (1) "you can get assistance from colleagues if you ask for it" and (2) "you can get assistance from your superior/boss if you ask for it." The index is rescaled to range between 0-100 (higher score indicates a higher degree of relatedness)</p>
Log monthly income	<p>Log monthly income in Euros, and PPP-adjusted using Purchasing power parities (PPPs), price level indices and real expenditures for ESA 2010 aggregates, actual individual consumption from Eurostat. The income information for 2005 is based on taking the midpoint of the country-specific income intervals and then converting them to Euros. In 2010 and 2015, respondents reported their actual income amounts. Those who refused to do so were prompted to indicate their income on an interval. For 2010 when income was missing, we took the midpoint of all income intervals and added this information to the continuous income variable. In 2015, when income was missing, we added the median of each income interval based on the 1991-2015 cumulative file.</p> <p>The variable is coded as 1 if the respondent reported receiving any of the following: performance pay, profit sharing, income from company shares, advantages such as medical services, access to shops. The variable is coded as 0 if the respondent receives none of these benefits. 0= no fringe benefits/bonus; 1=some fringe benefits/bonus.</p>
Benefits and performance pay	<p>A binary indicator variable based on the statement "My job offers good prospects for career advancement." The response has been recoded from the original agree-disagree scale, whereby 1 denotes if the respondent strongly agrees or agrees with the statement and 0 if they are neutral, disagree, or strongly disagree with the statement. 0=no career advancement opportunities; 1=career advancement opportunities.</p>
Career advancement	

Continued on next page

Table 4.A.2 – Continued from previous page

Variables	Explanation
Job insecurity	A binary indicator variable based on the statement “I might lose my job in the next 6 months.” The response has been recorded from the original agree-disagree scale, whereby 1 denotes if the respondent strongly agrees or agrees with the statement and 0 if they are neutral, disagree, or strongly disagree with the statement. 0=will not lose job; 1=may lose job in the next six months.
Log weekly work hours	Log of usual weekly hours worked per week in the main job.
Control variables	Age (in years); male (0 = female and 1 = male); household size (number of people in household); spouse in household (1=spouse/partner; 0=no spouse/partner); presence of children in the household (1=yes, 0=no); education (1 = primary education or less (no education, early childhood education and primary education); 2= secondary (lower secondary education and upper secondary education), 3=tertiary (post-secondary non-tertiary education, short cycle tertiary education, bachelor or equivalent, master or equivalent, and doctorate or equivalent), occupation dummies (ISCO 88 one-digit categories); industry of employer dummies; permanent contract (1=yes, 0=no); log of number of people supervised at work; company size indicators; public employee indicator (1=public employee; 0=not a public employee); number of days worked per week; work tenure (number of years with the current company); whether the respondent has other jobs (1=yes, 0=no); whether the respondent is involved in voluntary or charitable activity (1=yes, 0=no); interview duration (in minutes), number of people present during the interview; interview month; interview day; interviewer fixed effects.

Table 4.A.3: Summary statistics, selected variables

Variable	Mean	Std. Dev.
Meaningful work index	81.422	20.291
Autonomy	64.044	28.517
Competence	57.297	26.014
Relatedness	73.804	24.477
Monthly income (in Euros, PPP-adjusted)	3,373.150	19,270.250
Benefits and performance pay	0.302	0.459
Job insecurity	0.181	0.385
Career advancement	0.335	0.472
Log weekly hours	38.265	10.391
Age	41.026	11.286
Male	0.481	0.500
Primary education or less	0.149	0.357
Secondary education	0.711	0.453
Tertiary education	0.140	0.347

Source: Authors based on the European Working Conditions Surveys (EWCS) 2005-2015.

Notes: See Table 4.A.2 for variable definitions. N=48,420. Income and weekly hours are logged in the regression analyses.

Table 4.A.4: Robustness check with missing values indicator

	(1)
Autonomy	0.131*** (0.003)
Competence	0.043*** (0.003)
Relatedness	0.169*** (0.004)
<i>Income: Ref: bottom 25%</i>	
Q2	0.721*** (0.222)
Q3	0.567** (0.242)
Top 25 %	0.791*** (0.279)
Income quartile missing	1.175*** (0.308)
Benefits and performance pay = Yes	-0.147 (0.166)
Benefits and performance pay = Missing information	0.113 (0.750)
Job insecurity = Yes	-3.538*** (0.212)
Job insecurity = Missing information	-0.804*** (0.311)
Career advancement = Yes	4.814*** (0.148)
Career advancement = Missing information	2.409*** (0.452)
<i>Weekly hours: Ref: bottom 25%</i>	
Q2	-0.827*** (0.211)
Q3	-1.618*** (0.245)
Top 25 %	-1.958*** (0.225)
Weekly hours missing	-0.183 (0.653)
N	75,250
Adj.R ²	0.426

Source: Authors based on the European Working Conditions Surveys (EWCS) 2005-2015.

Notes: Robust standard errors in parentheses. The dependent variable is perceptions of being engaged in meaningful work, which is an index measured on a scale of 0 to 100. See Table 4.A.2 for variable definitions. All regressions include country and year fixed effects, interview controls (duration, number of people present during interview, interview month, and interview day, interviewer fixed effects), individual controls, and occupation and industry fixed effects. To prevent loss of information, all control variables except autonomy, relatedness, and competence, include a missing values indicator. This indicator has no economically meaningful interpretation. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table 4.A.5: Correlations table, meaningful work and other subjective well-being measures

	Meaningful work index	Satisfaction with working conditions	Stress at work ^{††}	Health satisfaction [†]	Work engagement [†]	Job enthusiasm [†]
Meaningful work index	1					
Satisf. with working conditions	0.328*	1				
Stress at work ^{††}	-0.081*	-0.218*	1			
Health satisfaction [†]	0.121*	0.276*	-0.109*	1		
Work engagement [†]	0.340*	0.251*	-0.042*	0.112*	1	
Job enthusiasm [†]	0.443*	0.436*	-0.113*	0.189*	0.426*	1

Source: Authors based on the European Working Conditions Surveys (EWCS) 2005-2015.

Notes: [†] variable only available in the 2015 survey.^{††} variable available in 2010 and 2015. Meaningful work index measures perceptions of being engaged in meaningful work, which is an index measured on a scale of 0 to 100. See Table 4.A.2 for variable definition. Satisfaction with working conditions is measured on a scale of 1 (completely dissatisfied) to 4 (completely satisfied). Stress at work is measured on a frequency scale, whereby 1=Never, 2=Rarely, 3=Sometimes, 4=Most of the time, 5=Always. Health satisfaction is measured on a scale of 1=very bad to 5=very good (with the middle category being 'fair'). Work engagement is based on answers to the statement "Time flies when I am working," whereby 1=Never, 2=Rarely, 3=Sometimes, 4=Most of the time, 5=Always. Job enthusiasm is based on the statement "I am enthusiastic about my job," whereby 1=Never, 2=Rarely, 3=Sometimes, 4=Most of the time, 5=Always. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Bibliography

- Abowd, J. M. and M. H. Stinson (2013): "Estimating Measurement Error in Annual Job Earnings: A Comparison of Survey and Administrative Data," *The Review of Economics and Statistics*, 95, 1451–1467.
- Acemoglu, D. and D. H. Autor (2011): "Skills, Tasks and Technologies: Implications for Employment and Earnings," *Handbook of Labor Economics*, 4, 1043–1171.
- Acton, R. K. (2021): "Community College Program Choices in the Wake of Local Job Losses," *Journal of Labor Economics*, 39, 1129–1154.
- Akçomak, S., S. Kok, and H. Rojas-Romagosa (2016): "Technology, offshoring and the task content of occupations in the United Kingdom," *International Labour Review*, 155, 201–230.
- Aleksynska, M. (2018): "Temporary employment, work quality, and job satisfaction," *Journal of Comparative Economics*, 46, 722–735.
- Allan, B. A., C. Batz-Barbarich, H. M. Sterling, and L. Tay (2019): "Outcomes of Meaningful Work: A Meta-Analysis," *Journal of Management Studies*, 56, 500–528.
- Altonji, J. G., E. Blom, and C. Meghir (2012): "Heterogeneity in Human Capital Investments: High School Curriculum, College Major, and Careers," *Annual Review of Economics*, 4, 185–223, publisher: Annual Reviews.
- Angrave, D. and A. Charlwood (2015): "What is the Relationship Between Long Working Hours, Over-employment, Under-employment and the Subjective Well-being of Workers? Longitudinal Evidence From the UK," *Human Relations*, 68, 1491–1515.
- Arcidiacono, P. (2004): "Ability sorting and the returns to college major," *Journal of Econometrics*, 121, 343–375.
- Ariely, D., E. Kamenica, and D. Prelec (2008): "Man's Search for Meaning: The Case of Legos," *Journal of Economic Behavior & Organization*, 67, 671–677.
- Arntz, M., T. Gregory, and U. Zierahn (2017): "Revisiting the risk of automation," *Economics Letters*, 159, 157–160.
- Ashford, S. J., C. Lee, and P. Bobko (1989): "Content, Cause, and Consequences of Job Insecurity: A Theory-Based Measure and Substantive Test." *Academy of Management Journal*, 32, 803–829.
- Autor, D. and A. Salomons (2017): "Robocalypse Now—Does Productivity Growth Threaten Employment?" in *ECB Forum on Central Banking Proceedings*.

- Autor, D. H. (2015): "Why Are There Still So Many Jobs? The History and Future of Workplace Automation," *Journal of Economic Perspectives*, 29, 3–30.
- Autor, D. H. and D. Dorn (2013): "The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market," *American Economic Review*, 103, 1553–1597.
- Autor, D. H., D. Dorn, and G. H. Hanson (2013): "The China Syndrome: Local Labor Market Effects of Import Competition in the United States," *American Economic Review*, 103, 2121–2168.
- Autor, D. H. and M. J. Handel (2013): "Putting Tasks to the Test: Human Capital, Job Tasks, and Wages," *Journal of Labor Economics*, 31, S59–S96.
- Autor, D. H., L. F. Katz, and M. S. Kearney (2006): "The Polarization of the U.S. Labor Market," *American Economic Review*, 96, 189–194.
- 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, 1279–1333.
- Bailey, C., M. Lips-Wiersma, A. Madden, R. Yeoman, M. Thompson, and N. Chalofsky (2019a): "The five Paradoxes of meaningful Work: Introduction to the special issue Meaningful Work: Prospects for the 21st Century," *Journal of Management Studies*, 56, 481–499.
- Bailey, C., R. Yeoman, A. Madden, M. Thompson, and G. Kerridge (2019b): "A review of the empirical literature on meaningful work: Progress and research agenda," *Human Resource Development Review*, 18, 83–113.
- Bäker, A. and M. Mechtel (2018): "The role of task meaning on output in groups: Experimental evidence," *Managerial and Decision Economics*, 39, 131–141.
- Bandura, A. (1997): *Self-Efficacy: The Exercise of Control*, New York: Worth Publishers, 1st edition ed.
- Barrett, M., E. Oborn, W. J. Orlikowski, and J. Yates (2011): "Reconfiguring Boundary Relations: Robotic Innovations in Pharmacy Work," *Organization Science*, 23, 1448–1466.
- Baumgarten, D., I. Geishecker, and H. Görg (2013): "Offshoring, tasks, and the skill-wage pattern," *European Economic Review*, 61, 132–152.

- Beaudry, P. and D. A. Green (2003): "Wages and Employment in the United States and Germany: What Explains the Differences?" *American Economic Review*, 93, 573–602.
- Benz, M. and B. S. Frey (2008): "Being independent is a great thing: Subjective evaluations of self-employment and hierarchy," *Economica*, 75, 362–383.
- Van den Berge, W. and B. Ter Weel (2015): "Middensegment onder druk: nieuwe kansen door technologie," CPB Policy Brief, CPB.
- Binder, M. (2016): "'Do it with joy!' Subjective well-being outcomes of working in non-profit organizations," *Journal of Economic Psychology*, 54, 64–84.
- Binder, M. and A. Coad (2013): "Life satisfaction and self-employment: A matching approach," *Small Business Economics*, 40, 1009–1033.
- Bisschop, P., J. Zwetsloot, B. Ter Weel, and J. Van Kesteren (2020): "De overgang van het mbo naar de arbeidsmarkt de positie van jongeren met een migratieachtergrond belicht," SEO-rapport 2020-05.
- Blanchflower, D. G. and A. J. Oswald (1998): "What makes an entrepreneur?" *Journal of Labor Economics*, 16, 26–60.
- Böckerman, P. and P. Ilmakunnas (2009): "Job Disamenities, Job Satisfaction, Quit Intentions, and Actual Separations: Putting the Pieces Together," *Industrial Relations: A Journal of Economy and Society*, 48, 73–96.
- (2012): "The job satisfaction-productivity nexus: A study using matched survey and register data," *ILR Review*, 65, 244–262.
- Bond, T. N. and K. Lang (2018): "The Sad Truth about Happiness Scales," *Journal of Political Economy*, 127, 1629–1640.
- Borghans, L., B. Ter Weel, and B. A. Weinberg (2014): "People Skills and the Labor-Market Outcomes of Underrepresented Groups," *ILR Review*, 67, 287–334.
- Borghans, L. and B. ter Weel (2007): "The diffusion of computers and the distribution of wages," *European Economic Review*, 51, 715–748.
- Van den Bosch, A., B. Busser, S. Canisius, and W. Daelemans (2007): "An efficient memory-based morphosyntactic tagger and parser for Dutch," *Computational Linguistics in the Netherlands: Selected Papers from the Seventeenth CLIN Meeting*, 99–114.

- Both-Nwabuwe, J. M. C., M. T. M. Dijkstra, and B. Beersma (2017): "Sweeping the Floor or Putting a Man on the Moon: How to Define and Measure Meaningful Work," *Frontiers in Psychology*, 8, 1–14.
- Bradler, C., R. Dur, S. Neckermann, and A. Non (2016): "Employee recognition and performance: A field experiment," *Management Science*, 62, 3085–3099.
- Bresnahan, T. F., E. Brynjolfsson, and L. M. Hitt (2002): "Information Technology, Workplace Organization, and the Demand for Skilled Labor: Firm-Level Evidence," *The Quarterly Journal of Economics*, 117, 339–376.
- Brown, A., A. Charlwood, C. Forde, and D. Spencer (2007): "Job quality and the economics of New Labour: A critical appraisal using subjective survey data," *Cambridge Journal of Economics*, 31, 941–971.
- Brown, A., A. Charlwood, and D. A. Spencer (2012): "Not all that it might seem: why job satisfaction is worth studying despite it being a poor summary measure of job quality," *Work, employment and society*, 26, 1007–1018.
- Bryce, A. (2018): "Finding meaning through work: Eudaimonic well-being and job type in the US and UK," Sheffield Economic Series No2018004.
- Brynjolfsson, E. and A. McAfee (2014): *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*, New York and London: Norton & Company.
- Budd, J. W. and D. A. Spencer (2015): "Worker well-being and the importance of work: Bridging the gap," *European Journal of Industrial Relations*, 21, 181–196.
- Bárány, Z. L. and C. Siegel (2018): "Job Polarization and Structural Change," *American Economic Journal: Macroeconomics*, 10, 57–89.
- (2020): "Biased technological change and employment reallocation," *Labour Economics*, 67, 101930.
- Böhm, M. J. (2020): "The price of polarization: Estimating task prices under routine-biased technical change," *Quantitative Economics*, 11, 761–799.
- Card, D., T. Lemieux, and W. C. Riddell (2018): "Unions and Wage Inequality: The Roles of Gender, Skill and Public Sector Employment," Working Paper 25313, National Bureau of Economic Research, series: Working Paper Series.
- Caroli, E. and M. Godard (2016): "Does job insecurity deteriorate health?" *Health economics*, 25, 131–147.

- Caroli, E. and J. Van Reenen (2001): "Skill-Biased Organizational Change? Evidence from A Panel of British and French Establishments," *The Quarterly Journal of Economics*, 116, 1449–1492.
- Cassar, L. and S. Meier (2018): "Nonmonetary incentives and the implications of work as a source of meaning," *Journal of Economic Perspectives*, 32, 215–238.
- Cassidy, H. (2017): "Task Variation Within Occupations," *Industrial Relations: A Journal of Economy and Society*, 56, 393–410.
- Cavaglia, C. and B. Etheridge (2017): "Job polarization, task prices and the distribution of task returns," Working Paper 2017-09, ISER Working Paper Series.
- Centraal Bureau voor de Statistiek (2016): "Helpt leert door na beroepsopleidende leerweg mbo," .
- (2018): "Minder mbo'ers naar hbo in afgelopen decennium," .
- Chadi, A. and C. Hetschko (2018): "The magic of the new: How job changes affect job satisfaction," *Journal of Economics & Management Strategy*, 27, 23–39.
- Chadi, A., S. Jeworrek, and V. Mertins (2016): "When the meaning of work has disappeared: Experimental evidence on employees' performance and emotions," *Management Science*, 63, 1696–1707.
- Chandler, D. and A. Kapelner (2013): "Breaking monotony with meaning: Motivation in crowdsourcing markets," *Journal of Economic Behavior & Organization*, 90, 123–133.
- Chater, N. and G. Loewenstein (2016): "The under-appreciated drive for sense-making," *Journal of Economic Behavior & Organization*, 126, 137–154.
- Clark, A. E. (2001): "What really matters in a job? Hedonic measurement using quit data," *Labour Economics*, 8, 223–242.
- (2005): "What Makes a Good Job? Evidence from OECD Countries," in *Job Quality and Employer Behavior* ., ed. by C. L. S. Bazen and W. Salverda, London: Palgrave Macmillan UK, 11–30.
- (2011): "Worker well-being in booms and busts," in *The Labour Market in Winter: The State of Working Britain.*, ed. by G. a. J. Wadsworth, Oxford and New York: Oxford University Press, 128–143.

- (2015): “What makes a good job? Job quality and job satisfaction,” *IZA World Of Labor*, 215, 1–10.
- Cnossen, F., M. Piracha, and G. Tchuente (2021): “Learning the Right Skill: The Returns to Social, Technical and Basic Skills for Middle-Educated Graduates,” GLO Discussion Paper No. 979, Global Labor Organization (GLO).
- Conti, G. and S. Pudney (2011): “Survey design and the analysis of satisfaction,” *Review of Economics and Statistics*, 93, 1087–1093.
- Cornelißen, T. (2009): “The interaction of job satisfaction, job search, and job changes: An empirical investigation with German panel data,” *Journal of Happiness Studies*, 10, 367–384.
- Cortes, G. M. (2015): “Where Have the Middle-Wage Workers Gone? A Study of Polarization Using Panel Data,” *Journal of Labor Economics*, 34, 63–105.
- Cottini, E. and C. Lucifora (2013): “Mental health and working conditions in European countries,” *Cornell University ILR Review*, 66, 958–988.
- Cuyper, N. D. and H. D. Witte (2006): “The impact of job insecurity and contract type on attitudes, well-being and behavioural reports: A psychological contract perspective,” *Journal of Occupational and Organizational Psychology*, 79, 395–409.
- D’Ambrosio, C., A. E. Clark, and M. Barazzetta (2018): “Unfairness at work: Well-being and quits,” *Labour Economics*, 51, 307–316.
- Dauth, W., S. Findeisen, J. Suedekum, and N. Woessner (2021): “The Adjustment of Labor Markets to Robots,” *Journal of the European Economic Association*.
- De Bustillo, R. M., E. Fernández-Macías, F. Esteve, and J.-I. Antón (2011a): “E pluribus unum? A critical survey of job quality indicators,” *Socio-Economic Review*, 9, 447–475.
- De Bustillo, R. M., E. Fernández-Macías, J.-I. Antón, and F. Esteve (2011b): *Measuring more than money: The social economics of job quality*, Cheltenham: Edward Elgar Publishing.
- De Witte, H., T. Vander Elst, and N. De Cuyper (2015): “Job Insecurity, Health and Well-Being,” in *Sustainable Working Lives: Managing Work Transitions and Health throughout the Life Course*, ed. by J. Vuori, R. Blonk, and R. H. Price, Dordrecht: Springer Netherlands, Aligning Perspectives on Health, Safety and Well-Being, 109–128.

- Deaton, A. (1985): "Panel data from time series of cross sections," *Journal of Econometrics*, 30, 109–126.
- Deci, E. and R. Ryan (2008): "Hedonia, eudaimonia, and well-being: An introduction," *Journal of Happiness Studies*, 9, 1–11.
- Deci, E. and R. M. Ryan (1985): *Intrinsic motivation and self-determination in human behavior*, New York: Plenum Press.
- Deming, D. (2017): "The Growing Importance of Social Skills in the Labor Market," *The Quarterly Journal of Economics*, 132, 1593–1640.
- Deming, D. and L. B. Kahn (2017): "Skill Requirements across Firms and Labor Markets: Evidence from Job Postings for Professionals," *Journal of Labor Economics*, 36, S337–S369.
- Deming, D. and K. L. Noray (2018): "STEM Careers and the Changing Skill Requirements of Work," Working Paper 25065, National Bureau of Economic Research.
- Deming, D. J. (2021): "The Growing Importance of Decision-Making on the Job," Working Paper 28733, National Bureau of Economic Research.
- Dur, R. and M. van Lent (2019): "Socially useless jobs," *Industrial Relations: A Journal of Economy and Society*, 58, 3–16.
- Edin, P.-A., T. Evans, G. Graetz, S. Hernnäs, and G. Michaels (2019): "Individual Consequences of Occupational Decline," Tech. Rep. CEPDP1629, Centre for Economic Performance, LSE.
- Eeckhout, J. and P. Kircher (2011): "Identifying Sorting—In Theory," *The Review of Economic Studies*, 78, 872–906.
- Eggenberger, C., M. Rinawi, and U. Backes-Gellner (2018): "Occupational specificity: A new measurement based on training curricula and its effect on labor market outcomes," *Labour Economics*, 51, 97–107.
- Entorf, H., M. Gollac, and F. Kramarz (1999): "New Technologies, Wages, and Worker Selection," *Journal of Labor Economics*, 17, 464–491.
- Eurofound (2012): *Trends in job quality in Europe*, Luxembourg: Publications Office of the European Union.
- Eurofound (2019): "European Working Conditions Survey Integrated Data File, 1991–2015. [Data Collection] 7th edition," Tech. rep., UK Data Service.

- European Foundation for the Improvement of Living and Working Conditions (2007): "European Working Conditions Survey 2005 [data collection]. UK Data Service. SN: 5639," .
- (2012): "European Working Conditions Survey 2010 [data collection]. UK Data Service," .
- (2017): "European Working Conditions Survey 2015 [data collection]. 4th Edition. UK Data Service. SN: 8098," .
- (2019): "European Working Conditions Survey Integrated Data File, 1991-2015. [data collection]. 7th Edition. UK Data Service. SN: 7363," .
- Fehrler, S. and M. Kosfeld (2014): "Pro-social missions and worker motivation: An experimental study," *Journal of Economic Behavior & Organization*, 100, 99–110.
- Felstead, A., D. Gallie, F. Green, and G. Henseke (2019): "Conceiving, designing and trailing a short-form measure of job quality: A proof-of-concept study," *Industrial Relations Journal*, 50, 2–19.
- Felstead, A., D. Gallie, F. Green, and Y. Zhou (2007): "Skills At Work, 1986 to 2006," Tech. rep.
- Fouarge, D., A. Künn, and D. Punt (2017): "De rol van arbeidsmarktinformatie in de opleidingskeuze van mbo'ers," ROA Reports No. 009.
- 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.
- Gentzkow, M., B. Kelly, and M. Taddy (2019): "Text as Data," *Journal of Economic Literature*, 57, 535–574.
- Georgellis, Y. and A. Yusuf (2016): "Is becoming self-employed a panacea for job satisfaction? Longitudinal evidence from work to self-employment transitions," *Journal of Small Business Management*, 1, 53–76.
- Goldin, C. and L. F. Katz (2010): *The Race between Education and Technology*, Cambridge, Massachusetts; London, England: Belknap Press.
- Golsteyn, B. H. H. and A. Stenberg (2017): "Earnings over the Life Course: General versus Vocational Education," *Journal of Human Capital*, 11, 167–212.

- Goos, M. and A. Manning (2007): "Lousy and Lovely Jobs: The Rising Polarization of Work in Britain," *The Review of Economics and Statistics*, 89, 118–133.
- Goos, M., A. Manning, and A. Salomons (2009): "Job Polarization in Europe," *American Economic Review*, 99, 58–63.
- (2014): "Explaining Job Polarization: Routine-Biased Technological Change and Offshoring," *American Economic Review*, 104, 2509–2526.
- Goos, M. and A. Salomons (2009): "De veranderende structuur en kwaliteit van banen in België," *Over.werk. Tijdschrift van het Steunpunt WSE*, 19, 30–42.
- Graham, C., K. Laffan, and S. Pinto (2018): "Well-being in metrics and policy," *Science*, 362, 287–288.
- Graham, C. and M. Nikolova (2015): "Bentham or Aristotle in the development process? An empirical investigation of capabilities and subjective well-being," *World development*, 68, 163–179.
- Grant, A. M. (2008): "The significance of task significance: Job performance effects, relational mechanisms, and boundary conditions," *Journal of Applied Psychology*, 93, 108.
- Green, F. (2006): *Demanding Work: The Paradox of Job Quality in the Affluent Economy*, Princeton: Princeton University Press.
- (2010): "Well-being, job satisfaction and labour mobility," *Labour Economics*, 17, 897–903.
- Green, F., A. Felstead, D. Gallie, and H. Inanc (2016): "Job-related well-being through the Great Recession," *Journal of Happiness Studies*, 17, 389–411.
- Greenwood, J., N. Guner, G. Kocharkov, and C. Santos (2016): "Technology and the Changing Family: A Unified Model of Marriage, Divorce, Educational Attainment, and Married Female Labor-Force Participation," *American Economic Journal: Macroeconomics*, 8, 1–41.
- Gregory, T., A. Salomons, and U. Zierahn (2021): "Racing with or Against the Machine? Evidence from Europe," *Journal of the European Economic Association*.
- Grossman, G. M. and E. Rossi-Hansberg (2008): "Trading Tasks: A Simple Theory of Offshoring," *American Economic Review*, 98, 1978–1997.

- Haller, S. A. and I. Siedschlag (2011): "Determinants of ICT adoption: evidence from firm-level data," *Applied Economics*, 43, 3775–3788.
- Hamermesh, D. S. (2018): "Why are professors 'Poorly paid'?" *Economics of Education Review*, 66, 137–141.
- Hamilton, B. H. (2000): "Does entrepreneurship pay? An empirical analysis of the returns to self-employment," *Journal of Political Economy*, 108, 604–631.
- Hansen, S., T. Ramdas, R. Sadun, and J. Fuller (2021): "The Demand for Executive Skills," Working Paper 28959, National Bureau of Economic Research.
- Hanushek, E. A., G. Schwerdt, L. Woessmann, and L. Zhang (2017): "General Education, Vocational Education, and Labor-Market Outcomes over the Lifecycle," *Journal of Human Resources*, 52, 48–87.
- Hessels, J., E. Arampatzi, P. van der Zwan, and M. Burger (2018): "Life satisfaction and self-employment in different types of occupations," *Applied Economics Letters*, 25, 734–740.
- Hooftman, W., G. Mars, B. Janssen, E. d. Vroome, M. Ramaekers, and S. v. d. Bossche (2018): *Nationale Enquête Arbeidsomstandigheden 2017. Methodologie en globale resultaten*, TNO ; CBS.
- Howell, D. R. and A. L. Kalleberg (2019): "Declining job quality in the United States: Explanations and evidence," *RSF: The Russell Sage Foundation Journal of the Social Sciences*, 5, 1–53.
- Hunt, J. and R. Nunn (2019): "Is Employment Polarization Informative About Wage Inequality and Is Employment Really Polarizing?" Working Paper 26064, National Bureau of Economic Research.
- Hyttinen, A., P. Ilmakunnas, and O. Toivanen (2013): "The return-to-entrepreneurship puzzle," *Labour Economics*, 20, 57–67.
- Israeli, O. (2007): "A Shapley-based decomposition of the R-square of a linear regression," *The Journal of Economic Inequality*, 5, 199–212.
- Jovanovic, B. and P. Rousseau (2005): "General Purpose Technologies," *Handbook of Economic Growth*, Elsevier.
- Kaplan, G. and S. Schulhofer-Wohl (2018): "The changing (dis-)utility of work," *Journal of Economic Perspectives*, 32, 239–258.

- Karlsson, N., G. Loewenstein, and J. McCafferty (2004): "The economics of meaning," *Nordic Journal of Political Economy*, 30, 61–75.
- Kassenboehmer, S. C. and J. P. Haisken-DeNew (2009): "You're fired! The causal negative effect of entry unemployment on life satisfaction," *The Economic Journal*, 119, 448–462.
- Katz, L. F. and K. M. Murphy (1992): "Changes in Relative Wages, 1963-1987: Supply and Demand Factors," *The Quarterly Journal of Economics*, 107, 35–78.
- Knabe, A. and S. Rätzel (2011a): "Quantifying the psychological costs of unemployment: The role of permanent income," *Applied Economics*, 43, 2751–2763.
- (2011b): "Scarring or scaring? The psychological impact of past unemployment and future unemployment risk," *Economica*, 78, 283–293.
- Knox, A., C. Warhurst, D. Nickson, and E. Dutton (2015): "More than a feeling: Using hotel room attendants to improve understanding of job quality," *The International Journal of Human Resource Management*, 26, 1547–1567.
- de Koning, J., A. Gelderblom, L. Den Hartog, and T. Berretty (2010): "Techniek: Exact goed? Het keuzeproces van allochtone en autochtone leerlingen in het (v)mbo verklaard." Tech. rep., SEOR, Erasmus University Rotterdam, Rotterdam.
- Kosfeld, M., S. Neckermann, and X. Yang (2017): "The effects of financial and recognition incentives across work contexts: The role of meaning," *Economic Inquiry*, 55, 237–247.
- Kurer, T. and B. Palier (2019): "Shrinking and shouting: the political revolt of the declining middle in times of employment polarization," *Research & Politics*, 6.
- Lazear, E. P. (2018): "Compensation and incentives in the workplace," *Journal of Economic Perspectives*, 32, 195–214.
- Leschke, J. and A. Watt (2014): "Challenges in constructing a multi-dimensional European job quality index," *Social indicators research*, 118, 1–31.
- Lewandowski, P., A. Park, W. Hardy, and Y. Du (2019): "Technology, Skills, and Globalization: Explaining International Differences in Routine and Nonroutine Work Using Survey Data," Tech. Rep. 12339, Institute of Labor Economics (IZA).
- Lysova, E. I., B. A. Allan, B. J. Dik, R. D. Duffy, and M. F. Steger (2019): "Fostering meaningful work in organizations: A multi-level review and integration," *Journal of Vocational Behavior*, 110, 374–389.

- Lévy-Garboua, L., C. Montmarquette, and V. Simonnet (2007): "Job satisfaction and quits," *Labour Economics*, 14, 251–268.
- MacKerron, G. (2012): "Happiness economics from 35 000 feet," *Journal of Economic Surveys*, 26, 705–735.
- MacPhee, D., S. Farro, and S. S. Canetto (2013): "Academic Self-Efficacy and Performance of Underrepresented STEM Majors: Gender, Ethnic, and Social Class Patterns," *Analyses of Social Issues and Public Policy*, 13, 347–369.
- Maestas, N., K. J. Mullen, D. Powell, T. Von Wachter, and J. B. Wenger (2017): "Working Conditions in the United States Results of the 2015 American Working Conditions Survey," SSRN Scholarly Paper ID 3183866, Social Science Research Network, Rochester, NY.
- Malamud, O. and C. Pop-Eleches (2010): "General Education versus Vocational Training: Evidence from an Economy in Transition," *The Review of Economics and Statistics*, 92, 43–60.
- Marcolin, L., S. Miroudot, and M. Squicciarini (2016): "The Routine Content Of Occupations," OECD Trade Policy Papers, Organisation for Economic Co-operation and Development, Paris.
- Martela, F. and T. J. Riekkki (2018): "Autonomy, Competence, Relatedness, and Beneficence: A Multicultural Comparison of the Four Pathways to Meaningful Work," *Frontiers in Psychology*, 9, 1157.
- Martela, F., R. M. Ryan, and M. F. Steger (2018): "Meaningfulness as Satisfaction of Autonomy, Competence, Relatedness, and Beneficence: Comparing the Four Satisfactions and Positive Affect as Predictors of Meaning in Life," *Journal of happiness studies*, 19, 1261–1282.
- Merino-Salazar, P., L. Artazcoz, J. Campos-Serna, D. Gimeno, and F. G. Benavides (2015): "National working conditions surveys in Latin America: comparison of methodological characteristics," *International Journal of Occupational and Environmental Health*, 21, 266–274.
- Meyer, B. D. and N. Mittag (2019a): "Combining Administrative and Survey Data to Improve Income Measurement," Working Paper 25738, National Bureau of Economic Research.
- (2019b): "An Empirical Total Survey Error Decomposition Using Data Combination," Working Paper 25737, National Bureau of Economic Research.

- Michaels, G., F. Rauch, and S. J. Redding (2019): "Task Specialization in U.S. Cities from 1880 to 2000," *Journal of the European Economic Association*, 17, 754–798.
- Mincer, J. (1974): "Schooling, Experience, and Earnings," NBER Books, National Bureau of Economic Research, Inc.
- National Center for O*NET Development (2021): "Skills," .
- Nikolova, M. (2019): "Switching to self-employment can be good for your health," *Journal of Business Venturing*, 34, 664–691.
- Nikolova, M. and S. H. Ayhan (2019): "Your spouse is fired!" *How much do you care?* *Journal of Population Economics*, 32, 799–844.
- Nikolova, M. and F. Cnossen (2020): "What makes work meaningful and why economists should care about it," *Labour Economics*, 65, 101847.
- Nikolova, M. and C. Graham (2014): "Employment, Late-Life Work, Retirement, and Well-being in Europe and the U." *S. IZA Journal of European Labor Studies*, 3, 1–30.
- OECD (2013): *OECD guidelines on Measuring subjective well-being*, Organisation for Economic Cooperation and Development Publishing.
- Olsson, U. (1979a): "Maximum likelihood estimation of the polychoric correlation coefficient," *Psychometrika*, 44, 443–460.
- (1979b): "Maximum likelihood estimation of the polychoric correlation coefficient," *Psychometrika*, 44, 443–460.
- Ministerie van Onderwijs, C. e. W. (2013): "Home - Inspectorate of Education," .
- Parent-Thirion, A., I. Biletta, J. Cabrita, O. Llave Vargas, G. Vermeylen, A. Wilczynska, and M. Wilkens (2017): *6th European Working Conditions Survey: overview report*, no. 16/34 in EF, Luxembourg: Publications Office of the European Union, 2017 update ed., oCLC: 1019729874.
- Park, J. and N. Lee (2009): "First Korean Working Conditions Survey: a comparison between South Korea and EU countries," *Industrial Health*, 47, 50–54.
- Preston, A. E. (1989): "The nonprofit worker in a for-profit world," *Journal of Labor Economics*, 7, 438–463.
- Rehdanz, K. and D. Maddison (2005): "Climate and happiness," *Ecological Economics*, 52, 111–125.

- Reichert, A. R. and H. Tauchmann (2017): "Workforce reduction, subjective job insecurity, and mental health," *Journal of Economic Behavior & Organization*, 133, 187–212.
- Reijnders, L. S. M. and G. J. de Vries (2017): "Technology, Offshoring and the Rise of Non-Routine Jobs," GGDC Research Memorandum 167, Groningen.
- Reinhold, M. and S. Thomsen (2017): "The changing situation of labor market entrants in Germany," *Journal for Labour Market Research*, 50, 161–174.
- Rica, S. D. L., L. Gortazar, and P. Lewandowski (2020): "Job Tasks and Wages in Developed Countries: Evidence from PIAAC," *Labour Economics*, 101845.
- Ross, M. B. (2017): "Routine-biased technical change: Panel evidence of task orientation and wage effects," *Labour Economics*, 48, 198–214.
- Rosso, B. D., K. H. Dekas, and A. Wrzesniewski (2010): "On the meaning of work: A theoretical integration and review," *Research in Organizational Behavior*, 30, 91–127.
- Rothwell, J. and S. Crabtree (2019): Not just a job: New evidence on the quality of work in the United States.
- Roy, A. D. (1951): "Some Thoughts on the Distribution of Earnings," *Oxford Economic Papers, New Series*, 3, 135–146.
- Ryan, R. M. and E. L. Deci (2000a): "Self-determination theory and the facilitation of intrinsic motivation, social development, and well-being," *American Psychologist*, 55, 68–78.
- (2000b): "Self-determination theory and the facilitation of intrinsic motivation, social development, and well-being," *American Psychologist*, 55, 68–78.
- Ryan, R. M., V. Huta, and E. L. Deci (2008): "Living well: A self-determination theory perspective on eudaimonia," *Journal of happiness studies*, 9, 139–170.
- Salverda, W. and D. Checchi (2015): "Labor Market Institutions and the Dispersion of Wage Earnings (Chapter 18)," in *Handbook of Income Distribution*, ed. by A. B. Atkinson and F. Bourguignon, Elsevier, vol. 2, 1535–1727.
- SBB (2021): "About SBB," .
- Schakel, W. and D. Van der Pas (2021): "Degrees of influence: Educational inequality in policy representation," *European Journal of Political Research*, 60, 418–437.

- Sen, A. (1999): *Development as Freedom*, New York, NY: Anchor Books.
- Shiller, R. J. (2019): "Narratives about Technology-Induced Job Degradations Then and Now," Working Paper 25536, National Bureau of Economic Research, series: Working Paper Series.
- Shorrocks, A. F. (2013): "Decomposition procedures for distributional analysis: a unified framework based on the Shapley value," *Journal of Economic Inequality*, 1–28.
- Smids, J., S. Nyholm, and H. Berkers (2020): "Robots in the Workplace: a Threat to—or Opportunity for—Meaningful Work?" *Philosophy & Technology*, 33, 503–522.
- Sousa-Poza, A. and A. A. Sousa-Poza (2000): "Well-being at work: a cross-national analysis of the levels and determinants of job satisfaction," *The Journal of Socio-Economics*, 29, 517–538.
- Spencer, D. (2009): *The Political Economy of Work*, Abingdon and New York: Routledge.
- Spencer, D. A. (2015): "Developing an understanding of meaningful work in economics: the case for a heterodox economics of work," *Cambridge Journal of Economics*, 39, 675–688.
- Spitz-Oener, A. (2006): "Technical Change, Job Tasks, and Rising Educational Demands: Looking outside the Wage Structure," *Journal of Labor Economics*, 24, 235–270.
- Steger, M. F., B. J. Dik, and R. D. Duffy (2012): "Measuring Meaningful Work: The Work and Meaning Inventory (WAMI)," *Journal of Career Assessment*, 20, 322–337.
- Stern, S. (2004): "Do Scientists Pay to Be Scientists?" *Management Science*, 50, 835–853.
- Stinebrickner, R., T. Stinebrickner, and P. Sullivan (2018): "Job Tasks, Time Allocation, and Wages," *Journal of Labor Economics*, 37, 399–433.
- Stone, A. and C. Mackie (2014): *Subjective Well-Being: Measuring Happiness, Suffering, and Other Dimensions of Experience*: National Academies Press.
- Ter Weel, B. and S. Kok (2013): "De Nederlandse arbeidsmarkt in taken: Eerste bevindingen uit de Nederlandse Skills Survey," Tech. rep., CPB, Den Haag.
- Ter Weel, B., J. Zwetsloot, and P. Bisschop (2021): "Technologie verslechtert arbeidsmarktkansen van mbo'ers," *Economisch Statistische Berichten*.

- Terzidis, N. and R. Ortega-Argilés (2021): "Employment polarization in regional labor markets: Evidence from the Netherlands," *Journal of Regional Science*, 1–31.
- Tims, M. and A. B. Bakker (2010): "Job crafting: Towards a new model of individual job redesign," *SA Journal of Industrial Psychology*, 36, e1–e9.
- Vohs, K. D., N. L. Mead, and M. R. Goode (2006): "The psychological consequences of money," *Science*, 314, 1154–1156.
- Walk, M. and F. Handy (2018): "Job Crafting as Reaction to Organizational Change," *The Journal of Applied Behavioral Science*, 54, 349–370.
- Wanous, J. P., A. E. Reichers, and M. J. Hudy (1997): "Overall job satisfaction: How good are single-item measures?" *Journal of Applied Psychology*, 82, 247–252.
- Warr, P. (1990): "The measurement of well-being and other aspects of mental health," *Journal of occupational Psychology*, 63, 193–210.
- (2007): *Work, happiness, and unhappiness*, New York and Abingdon: Routledge.
- Webb, M. (2019): "The Impact of Artificial Intelligence on the Labor Market," *SSRN Electronic Journal*.
- Weil, D. (2014): *The Fissured Workplace: Why Work Became So Bad for So Many and What Can Be Done to Improve It*, Cambridge, Massachusetts: Harvard University Press.
- Weiss, H. M. (2002): "Deconstructing job satisfaction: Separating evaluations, beliefs and affective experiences," *Human resource management review*, 12, 173–194.
- Winkelmann, L. and R. Winkelmann (1998): "Why are the unemployed so unhappy? Evidence from panel data," *Economica*, 65, 1–15.
- Wolf, T., M. Metzger, and R. E. Lucas (2019): "Experienced well-being and labor market status: The role of pleasure and meaning," SOEP papers on Multidisciplinary Panel Data Research No. 1043.
- Wolfe, M. T. and P. C. Patel (2019): "Exploring the differences in perceptions of work importance and job usefulness to society between self-employed and employed individuals," *Journal of Business Venturing Insights*, 12.
- 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, 289–325.
- Wrzesniewski, A. and J. E. Dutton (2001): "Crafting a job: Revisioning employees as active crafters of their work," *Academy of Management Review*, 26, 179–201.

English summary

Technology changes the world around us. It is impossible to imagine an office without computers, healthcare without medical technology and factories without robots. This has heavily affected the organization of work. As our workplaces change, so do the tasks we execute at work - and the skills we need to do them. This bears consequences for our careers: people who can use technology to improve their productivity thrive on technological progress, but that is not the case for workers that execute tasks that computers are relatively good at. Over the past decades, mainly such routine-intensive tasks have been replaced by technologies: these are short and repetitive movements, where many machines hold the comparative advantage against humans. On the other hand, demand has increased for nonroutine tasks that require problem-solving, creativity and abstract thinking – where technology has not (yet) been very efficient in.

Both direct and indirect, the chapters in this dissertation are set against this background: they concern the well-being of workers in a labour market that can be characterized by continuous technological change. The first chapter concerns the changing demand for routine-intensive tasks and the consequences that has for the careers of Dutch employees. The second chapter involves the changing demand for skills. This is especially relevant for the middle-skilled segment of the labour market, that has come under pressure in the past decades due to technological change. Therefore, this chapter focuses on middle-educated students that are entering the labour market in the 2010's. What should they learn in order to ensure a good start on the labour market? The last chapter's theme is somewhat uncommon for economists. Whereas the first two chapters focused on the contents of our work (tasks and skills), this chapter deals with the question whether workers think of their work as useful and fulfilling. The focus shifts from material gains from labour, to the immaterial meaning we derive from work. The chapter shows that meaningful work is an important topic to study, also for economists – which possibly will become more so in the future, when technology is able to replace even more tasks.

In **Chapter 2**, I study the role of routine tasks in shaping careers of Dutch employees. Since computers hold a comparative advantage in executing repetitive, rules-based tasks, one can expect that workers that execute such tasks fare relatively poor on the labour market. Conversely, people that perform cognitive nonroutine tasks should be more successful. The results from this chapter confirm this: those that perform a relatively nonroutine set of tasks have higher wages, and faster wage growth over time. Increasing the set of routine tasks will not always relate to lower wages, but they will if these tasks are executed in computer-intensive industries. This is not surprising, as computers are generally more efficient in routine tasks.

The relationship between tasks and wages is widely documented in the literature, but the majority of this research uses occupation-level task data. This requires the implicit assumption that everyone in the same occupation executes the same tasks. However, that is not always the case, and technological change may exacerbate these within-occupation differences in tasks, if people adapt their tasks to a changing environment, while still be classified in the same occupation. One example of this adaptation is the changing tasks of cashiers since the introduction of self-checkout machines: the weight on social tasks increases, and the routine-intensive tasks (manually scanning products) decrease. This implies that occupation-level data may not always pick up the actual tasks people perform. The most important contribution of this chapter to the extant literature is that I am measuring tasks on the different and individual level, by using the Dutch Working Conditions Survey. I determine the relative (non-) routine-intensity of jobs for each of the 155.000 workers in the sample, by using questions about their work: are they required to solve problems independently? Are they allowed to change the sequence, speed and methods of their tasks? Do they perform a lot of repetitive tasks? By measuring tasks on the individual level, I can test the relation between tasks and wages in a more direct way. Like other work in this field, I confirm the existence of task-based inequalities, but I do so on a more fine-grained level than the literature before me.

By measuring task-based inequalities in wages we can paint a picture of the success of Dutch workers in their career. However, I adopt a rather economic view on the definition of a successful career: I observe the average wage that people earn and their wage growth in the years following the survey. To accompany these findings, a number of additional analyses show that workers with routine-intensive tasks are more often in a non-tenured contract, work fewer hours, and have a higher change of becoming unemployed. Moreover, they have lower job satisfaction, and – relevant during lockdowns – they have fewer possibilities to work from home. These non-pecuniary aspects of jobs can have a substantial impact on people's lives, through their influence on livelihoods and job dignity. This makes them important variables to take into account in this type of research. Technology not only creates inequalities in terms of wages, it also generates 'good' and 'bad' jobs in terms of broader job quality indicators.

Chapter 3 deals with the skills that people need in order to thrive in today's labour market. It is motivated by the empirical finding from the literature that the labour market opportunities of Dutch middle-educated graduates (mbo students) have decreased over the past years. This can roughly be attributed to two mechanisms. First, mbo'ers are relatively often trained for routine-intensive occupations. Second,

firms that invest in new technologies have a stronger tendency of retraining their current employees than to hire recent graduates – which especially hurts middle-educated students trained for routine professions. Another strand of literature focuses on the growing importance of certain skills: for example, social, cognitive analytical and technical skills appear more in vacancy-texts than a few years ago. This raises the question: what are the right skills for Dutch middle-educated students to learn to ensure a good start on the labour market? The main contribution of this chapter to the literature is the construction of a new data set on skills, plus the fact that I specifically focus on the skills of mbo-students. Even though the decline in the demand for middle-skilled workers is well-established in the literature, the majority of papers researching specific skills focus on high-skilled workers. This is both because of the complementarity between high-skilled workers and technology, and because of better availability of (vacancy) data for higher educated positions. As a result, we know relatively little about the skills that precisely the group of middle-educated students should learn to keep thriving in the labour market of the future.

To investigate this topic I constructed a new data set based on the curricula of Dutch mbo-degrees ("*kwalificatiedossiers*"). By means of Natural Language Programming techniques, I extracted verbs and nouns from skill descriptions in these text documents. From this list of verb-noun combinations I created a new index of skills that students learn, by labeling each of these combinations as either a social, technical or basic-cognitive skill. They look something like this: "discuss, colleague" is a social skill, "adjust, machine" is technical, and "read, document" is a basic-cognitive skill. In this way I can ascertain whether a certain degree is relatively intensive in social or technical skills. For instance, tourism programs are relatively more focused on social skills than administrative training, even though they are both in the same – Economic – field of education. I subsequently merge this new data to labour market data from all students that graduated in these degrees in the past ten years. In this way, I could analyse whether there exists a relationship between learning certain skills at school, and the wage in the first years after graduation.

The chapter presents three main findings. First, the more a degree focuses on technical skills, the higher the starting wage from their graduates is. This positive relation remains significant for at least ten years after graduating. Second, student from relatively social-intensive degrees have lower-than-average wages. This could indicate a lower demand for social skills for middle-educated students, but it might also be explained by the fact that social skills are best learned outside school, not in classrooms. This would also explain why the return is higher to apprenticeship-based

graduates, that have spent most of their education in the field. Third, the association between skills and wages differs significantly between levels and field of education. For example, the positive relation between technical skills and wages is specifically strong for graduates in the health sector. This might be indicative of complementarity between care tasks and knowing how to operate medical equipment.

The results from this chapter are still quite exploratory and further research is needed. It does seem clear that the search for new data sources can help to understand the demand for skills – with the final goal of being able to shape curricula based on empirical findings. I hope this chapter will stimulate further discussion and research on the topic, especially when middle-educated students seem to be most at risk for facing negative consequences of technological progress.

The last paper, presented in **Chapter 4**, deals with the concept of meaningful work. The academic discourse about this topic can be traced back as early as the economic science itself. Already Karl Marx reasoned that industrialization, the standardization of work and splitting-up of tasks, should lead to a decrease in meaning: employees would no longer have the idea their personal contributions matter to a final product. Later, economists added ‘compensating differentials’ to their labour supply models in order to compensate for meaningfulness: people are willing to earn less, as long as their tasks are perceived as meaningful. Modern economics pays relatively little attention to the concept. Yet, the interest seems to be increasing – possibly resulting from decreased bargaining powers of individual workers in the context of technological progress and globalization.

This chapter reinvigorates this interest, by combining insights from psychology with the economic labour supply model. We take the self-determination theory from Ryan and Deci as our starting point. In this theory, people will feel motivated when their psychological needs are met: they should experience autonomy, competence and relatedness at work. We show empirically that people with stronger experiences of these three elements also perceive their tasks as more fulfilling and useful – and thus meaningful. The three psychological needs are also stronger predictors of meaningful work than ‘hard’ variables, such as wages. Relevant to economists, this paper also shows that people that think of their work as more meaningful, also score higher on variables related to the supply of labour: they are less often ill, are more willing to participate in training and want to retire later.

Looking forward to the future of (the research on) work, I foresee a number of interesting directions. The ever growing availability of micro-level labour market data allows researchers to understand economic mechanisms on an increasingly more fundamental levels, and how large trends such as technological change and

globalization affect our work. New techniques in analysis, such as those related to Natural Language Processing, also provide new opportunities for generating data from sources that are unexploited by economists up until now. I also hope that the theme of dignifying work will receive more attention in the coming years: work is not solely a means to earn a living, and concepts as meaningful work, self-determination and job quality will likely become increasingly important, especially when technology will substitute more tasks. Being aware of these impacts now can also determine the direction in which technology will change: how technology-driven do we want our world to become, especially when it causes rising inequalities in both wages and job quality? Here lies a responsibility for policy makers and employers for making these choices, but there is most certainly also a responsibility for researchers to unveil the potential negative consequences of technology. This information is crucial in shaping policies, in order to ensure the future labour market can work for all of us.

Nederlandse samenvatting

De wereld om ons heen verandert door technologie. Het is moeilijk om je nog een kantoor zonder computers, een ziekenhuis zonder medische apparatuur of fabrieken zonder robots voor te stellen. Omdat de werkplaats waarin wij werken aan verandering onderhevig is, veranderen ook de taken die we uitvoeren op werk – en de vaardigheden die we daarvoor nodig hebben. Dit heeft consequenties voor onze carrières: mensen die door technologie hun productiviteit kunnen verbeteren varen goed op technologische vooruitgang, maar dat geldt niet voor de mensen die taken uitvoeren waar computers relatief efficiënt in zijn. Over de afgelopen decennia zijn voornamelijk dit soort routinematige taken vervangen door technologie: dit zijn korte, repetitieve werkzaamheden, waarin computers een comparatief voordeel hebben ten opzichte van mensen. Daarentegen steeg de vraag naar niet-routinematige taken, die vooral gaan over probleemoplossend vermogen, creativiteit en abstract denken – taken waar technologie tot nu nog niet heel goed in is gebleken.

De hoofdstukken in dit proefschrift spelen zich, direct of indirect, af tegen deze achtergrond: ze gaan over het welbevinden van werknemers op een arbeidsmarkt die gekarakteriseerd wordt door continue technologische verandering. Het eerste werk gaat over de veranderende vraag naar routinematige taken en de gevolgen daarvan voor de loopbanen van Nederlanders. Het tweede hoofdstuk gaat over de veranderende vraag naar vaardigheden. Dit is in het bijzonder relevant voor het middensegment van de arbeidsmarkt, dat de laatste jaren het meest onder druk is komen te staan door technologische verandering. Daarom focust dit hoofdstuk zich op middelbaar opgeleiden, die op dit moment de arbeidsmarkt betreden. Wat moeten we mbo-studenten leren om een goede start op de arbeidsmarkt te bewerkstelligen? Het laatste werk bekijkt technologische verandering vanuit een ander perspectief, dat wellicht wat ongewoon is voor een econoom. Waar in de eerste twee hoofdstukken centraal staat wat werknemers doen en kunnen, en welk loon daar tegenover staat, gaat dit hoofdstuk in op de vraag of werknemers hun werk nuttig vinden en het hun voldoening geeft. De focus verschuift daarmee van materiële verdiensten naar de immateriële betekenis die werk heeft. Ik laat zien dat de betekenis van werk ook voor economen het bestuderen waard is – en misschien juist als technologie nog meer taken kan vervangen in de toekomst.

In **Hoofdstuk 2** onderzoek ik of de routinematigheid van taken verschillen in de kwaliteit loopbanen van Nederlandse werknemers kan verklaren. Aangezien computers relatief beter zijn in het uitvoeren van repetitieve, gestandaardiseerde taken, is het de verwachting dat werknemers die dit soort taken uitvoeren het minder goed doen op de arbeidsmarkt. Omgekeerd zouden mensen die cognitieve niet-routinematige taken uitvoeren dan juist succesvoller moeten zijn. De resultaten laten

dit ook zien: mensen die relatief veel niet-routinematige taken uitvoeren hebben een hoger loon en sterkere loongroei over tijd. Meer routinematige taken uitvoeren gaat niet altijd gepaard met een lager loon, maar wel als deze taken worden uitgevoerd in computer-intensieve industrieën. Dat is niet verrassend, omdat computers over het algemeen efficiënter zijn in het uitvoeren van routinematige taken dan mensen.

Deze relatie tussen loon en taken is al vaker gedocumenteerd in de literatuur, maar het overgrote deel hiervan gebruikt daarvoor takendata op beroepsniveau. Dat betekent dat wordt aangenomen dat iedereen in hetzelfde beroep dezelfde taken uitvoert. Echter, dat is niet altijd het geval, en door technologische verandering is er potentieel zelfs nog meer variatie binnen beroepen omdat hun taken aanpassen aan hun veranderende omgeving – terwijl hun beroep hetzelfde blijft. Een voorbeeld hiervan is het veranderende takenpakket van kassamedewerkers door de introductie van zelfpinkassa's: er komt meer gewicht op sociale taken, en de routinematige taken (het scannen zelf) nemen af. Dit betekent dat beroependata niet altijd oppikt wat mensen daadwerkelijk uitvoeren op de werkvloer. De belangrijkste bijdrage van dit hoofdstuk aan de literatuur is dat ik taken op een andere manier meet, door gebruik te maken van de Nederlandse Enquête Arbeidsomstandigheden. Voor elk van de 155.000 werknemers in de steekproef bepaal ik de relatieve (niet-) routinematigheid van hun taken door vragen over hun werk: moeten ze problemen zelfstandig oplossen? Mogen ze de volgorde, snelheid en methode van hun werkzaamheden aanpassen? Moeten ze veel repetitieve taken uitvoeren? Door het relatieve belang van (niet-)routinematige taken op individueel in plaats van beroepsniveau te meten kan ik een directere analyse doen naar de relatie tussen taken en loon. Net als veel ander werk uit de literatuur bevestigt dit hoofdstuk het bestaan van taak-gebaseerde loonongelijkheid, maar het doet dit wel op een fijnmaziger niveau dan voorheen.

Door die taak-gebaseerde ongelijkheid in loon kunnen we een beeld schetsen van hoe succesvol Nederlanders zijn in hun loopbanen. De definitie van een succesvolle carrière in dit hoofdstuk is overigens vrij economisch: ik kijk naar het gemiddelde loon dat mensen verdienen en hun loongroei in de jaren na het invullen van de enquête. Echter, in een aantal aanvullende analyses laat ik ook zien dat mensen met routinematige taken ook kwalitatief minder succesvolle loopbanen kennen: ze hebben vaker een flexibel contract, werken minder uren, en ze hebben ook een grotere kans om werkloos te raken. Daarnaast geven zij aan minder tevreden te zijn met hun baan, en – relevant gedurende lockdowns – hebben zij minder mogelijkheden om hun werk vanuit huis uit te voeren. Ook deze niet-geldelijke maatstaven kunnen een grote impact hebben op het leven van werknemers, middels de invloed die ze hebben

op zekerheid van bestaan en waardigheid van werk. Het is daarom belangrijk ze mee te nemen in dit soort vraagstukken: technologie zorgt niet alleen voor ongelijkheid in salaris, het creëert ook in bredere zin een wig tussen goede banen en slechte banen.

Hoofdstuk 3 gaat over de vaardigheden die mensen nodig hebben in de arbeidsmarkt van de 21e eeuw. Het vertrekt vanuit de bevinding uit de literatuur dat de arbeidsmarktkansen van middelbaar opgeleiden in de afgelopen jaren zijn afgenomen. Dat heeft grofweg twee redenen: i) mbo'ers worden vaker opgeleid voor banen met relatief veel routinematige taken en ii) bedrijven die investeren in nieuwe technologie neigen er eerder naar hun huidige werknemers andere taken te geven, dan om te investeren in nieuwe jonge mensen – en al helemaal niet in jongeren die opgeleid zijn in routinematige taken. Andere vaardigheden worden daarentegen steeds belangrijker: sociale, cognitieve analytische en technische vaardigheden komen bijvoorbeeld steeds vaker terug in vacatures. Er is alleen nog geen duidelijk beeld over welke vaardigheden specifiek nodig lijken te zijn voor mbo-studenten. De grootste bijdrage van dit hoofdstuk is de ontwikkeling van een nieuwe dataset over skills en de focus op mbo-studenten. Want ook al is de daling in de vraag naar middelbare beroepen al uitgebreid gedocumenteerd, de meerderheid van de literatuur over specifieke vaardigheden focust zich nog op hoogopgeleide werknemers, vooral vanwege de complementariteit tussen hoogopgeleid werk en technologie en betere beschikbaarheid van (vacature)data voor hoogopgeleide functies. We weten daarom nog weinig over de vaardigheden die de middelbaar opgeleide groep werknemers *wel* zou moeten hebben om waardig werk te kunnen blijven hebben in de toekomst.

De data uit dit hoofdstuk is gebaseerd op de kwalificatiedossiers van alle Nederlandse mbo-opleidingen. Door het gebruik van tekstanalysetechnieken (*Natural Language Processing*) heb ik uit alle skill-beschrijvingen werkwoorden en zelfstandig naamwoorden getrokken. Uit deze lijst van werkwoord-zelfstandig naamwoord combinaties, die iets zeggen over de vaardigheid die studenten leren, heb ik een maatstaf van vaardigheden ontwikkeld. Elk van deze combinaties wordt gelabeld als ofwel een sociale vaardigheid, een technische of een basis-cognitieve vaardigheid. Dat ziet er ongeveer zo uit: “overleggen, collega” is een sociale vaardigheid, “afstellen, machine” is technisch, en “lezen, document” is een basis-cognitieve vaardigheid. Elke opleiding krijgt een score op basis van de relatieve frequentie van elk van de drie skills. Op deze manier wordt duidelijk hoe ‘sociaal’ of ‘technisch’ een bepaalde opleiding is. Zo is er binnen de toerisme-opleidingen meer aandacht voor sociale vaardigheden dan binnen administratieve opleidingen, hoewel die beide binnen dezelfde sector – Economie – vallen. Vervolgens heb ik deze nieuwe data over skills gekoppeld aan de arbeidsmarktgegevens van alle studenten die in de afgelopen tien

jaar zijn afgestudeerd in deze opleidingen, om te zien waar ze werken en hoeveel ze verdienen. Zo kon ik analyseren of er een relatie is tussen het leren van bepaalde skills, en het loon in de eerste jaren na afstuderen.

Het hoofdstuk presenteert drie centrale bevindingen. Ten eerste, hoe technischer een opleiding is, hoe hoger het startloon van afgestudeerden. Deze positieve relatie houdt stand gedurende minstens 10 jaar na afstuderen. Ten tweede, studenten uit relatief sociaal-intensieve opleidingen hebben juist een lager-dan-gemiddeld loon. Dit zou kunnen wijzen op een lage vraag naar sociale skills voor middelbaar opgeleiden, maar het kan ook verklaard worden door het feit dat je sociale vaardigheden beter leert in de praktijk dan op school. Dat zou ook verklaren waarom de opbrengst van sociale skills hoger is bij BBL-studenten – die hun opleiding voor het overgrote deel in de praktijk hebben genoten. Ten derde, de relatie tussen vaardigheden en loon verschilt aanzienlijk tussen niveaus en sectoren. Zo is de positieve relatie tussen technische vaardigheden en loon vooral sterk voor afgestudeerden in het zorgdomein. Dit kan duiden op een sterke complementariteit tussen zorgtaken en technische kennis van aanwezige apparatuur.

De resultaten uit dit paper zijn vooralsnog vrij exploratief en verder vervolgonderzoek is nodig. Het lijkt wel duidelijk dat de zoektocht naar nieuwe databronnen ons kan helpen in het beter begrijpen van de vraag naar vaardigheden – om uiteindelijk iets te kunnen zeggen over de vaardigheden die studenten zouden moeten leren op school. Ik hoop met dit hoofdstuk die discussie en het onderzoek ernaar verder aan te zwengelen.

Het laatste artikel, gepresenteerd in **Hoofdstuk 4** handelt over het concept betekenisvol werk. De discussie hierover gaat zover terug als de economische wetenschap en zelf. Zo relateerde Karl Marx industrialisatie, het standaardiseren van werk en het opknippen van taken, aan het verliezen van betekenis: werknemers hebben dan niet langer het gevoel zelf een nuttige bijdrage te leveren aan de uitkomsten van een productieproces. Latere economen voegden ‘compenserende loonverschillen’ toe aan hun arbeidsmodellen: mensen zijn bereid loon in te leveren als hun werk maar betekenisvol is. In de moderne economie is de interesse in betekenisvol werk niet heel groot. Toch, het lijkt groeiend – wat een gevolg kan zijn van afgenomen marktmacht van individuele werknemers in de context van technologische verandering en globalisering.

Dit hoofdstuk hernieuwt deze interesse, door inzichten uit de psychologie te koppelen aan het economische arbeidsaanbodmodel. Wij nemen daarbij de zelfbeschikkingstheorie van Ryan en Deci als startpunt. Mensen voelen zich gemotiveerd als zij het gevoel hebben zelf te beschikken over hun werk: omdat ze autonomie er-

varen, zich competent achten en zich verbonden voelen met hun omgeving. We laten empirisch zien dat mensen die dit ervaren, ook vaker ervaren dat hun werk nuttig is en voldoening geeft – en dus betekenisvol is. De zelfbeschikkingsvariabelen blijken ook beter in het voorspellen van betekenis in werk dan ‘harde’ karakteristieken van banen, zoals het salaris. Relevant voor economen, laten we ook zien dat mensen die hun werk betekenisvol vinden ook hoger scoren op economische variabelen rondom aanbod van arbeid: ze zijn minder vaak ziek, doen vaker aan bijscholing, en willen op latere leeftijd met pensioen.

Vooruitkijkend op de toekomst van het onderzoek naar werk, zie ik meerdere relevante richtingen. De steeds groeiende beschikbaarheid van fijnmazige microdata stelt onderzoekers in staat om op een steeds dieper niveau te begrijpen hoe de economie werkt, en hoe grote trends als technologische verandering en globalisering ons werk beïnvloeden. Analysetechnieken, zoals die uit de hoek van *Natural Language Processing*, bieden ook nieuwe mogelijkheden om meer data te kunnen genereren uit bronnen die tot recent niet door economen gebruikt werden. Daarnaast hoop ik dat het thema waardig werk ook meer aandacht zal krijgen: de arbeidsmarkt gaat over meer dan geld verdienen, en concepten als bestaanszekerheid, betekenisvol werk, en zelfbeschikking zullen waarschijnlijk alleen maar belangrijker worden, juist als technologie meer taken zal kunnen overnemen. Ons daar nu al bewust van zijn bepaalt ook de richting waarin technologie verandert, want daarin hebben we keuzes te maken. Hoe technologisch willen we dat onze wereld wordt, vooral als dat zorgt voor groeiende ongelijkheid in zowel loon als waardigheid van werk? Waar willen we naartoe? Hier ligt een verantwoordelijkheid voor beleidsmakers en het bedrijfsleven om dat soort keuzes te maken, maar die verantwoordelijk ligt ook zeker bij onderzoekers voor het aan het licht brengen van de potentiële negatieve consequenties van technologie – zodat we kunnen zorgen voor een arbeidsmarkt die voor iedereen werkt.

Acknowledgments

To a labour economist, writing a PhD dissertation can be seen as a process of applying skills to tasks, in order to achieve a meaningful outcome. As many who have gone before me will know, it felt as slightly more than just that. For the sake of avoiding cliches, I do not want to call it a journey - but it certainly was a trip. Here, I would like to thank those who have helped me along the way, and without whom this dissertation would not be the way it is now.

First, I would like to thank my supervisors: Robert Inklaar, Steven Brakman and Milena Nikolova. I feel very fortunate in having such a good team. Robert, voor je begeleiding op zowel werk- als persoonlijk vlak. Bedankt voor het mogelijk maken van mijn PhD, ondanks mijn twijfel in het begin. Ik heb altijd gevoeld dat je erin geloofde dat ik het kon afmaken, en hier staan we dan! Bedankt ook voor onze wekelijkse koffies op de maandagochtend: belangrijke maar vaak ook gewoon gezellige ijkpunten in de eentonigheid van de lockdownperiode. Steven, dank voor je vertrouwen in een goede afloop, ook al was het soms moeilijk om hoogte te krijgen van wat ik nou precies aan het doen was. Milena, I would not have been here without you: your continuous support, keen eye for details, introducing me to people, places and projects has given my PhD such an incredible boost. I am forever grateful for the hard work you have put into supervising this process.

I want to thank the members of the reading committee, Rob Alessie, Maarten Goos and Wolfgang Dauth, for reading my thesis and providing constructive feedback. Special thanks also to Maarten Goos for allowing me to take part in the Advanced Labour Economics course in 2019, which was very helpful in gaining better insights into the literature where this dissertation departs from.

I would also like to thank the people at SOM, Kristian, Arthur, Rina, Ellen en Hanneke, and my colleagues at the GEM department and the GGDC. Furthermore, I would like to thank the people at IZA. Even though my stay was unfortunately cut short due to Covid-19, I very much enjoyed being there – with specifically also very fond memories of my first time celebrating carnival. Thanks also to my new department, Economic Geography at the Faculty of Spatial Sciences, for your warm welcome and giving my all the space (!) in my first months here to finish the last parts of my PhD.

Next, a big thanks to my co-authors at the GLO, Matloob Piracha and Guy Tchuente. I am not sure whether you've realised how much of a help you two have been during the working-from-home period of my PhD. Our meetings really pulled me through, and provided some much needed perspective and motivation to continue working on the project - besides being good fun at times too! In this, I would also like to thank Olena Nizalova and Klaus Zimmermann for providing me

the opportunity to take part in the GLO Virtual Young Scholar Program, it was a great way to meet young scholars all over the globe and share our work in times of limited human interaction.

I have had some amazing teaching experiences in the past years. Mariko, thanks for being the best coordinator of a course one could wish for – I have very much enjoyed teaching IBE for these years. I would also like to thank the students from Rebuilding Education for your inspiring ideas and energy - despite having to either switch to or do everything online. And Jorim, for being a great teaching partner at the Rebuilding Education Leadership Lab.

Thanks also to my fellow PhD students for the good times: the drinks, the laughs, the lunches and the interesting conversations I have had with many of you. I would like to thank a few of my PhD buddies in particular. First, my two paranymphs: there are few people I feel so comfortable and completely myself with as with you both. Jos, my first roommate, we've had some great times in room 915 (I still have the office name card!). I admire your kindheartedness, honesty and how happy you can be for others. Juliette, whom I consider my roommate on many Tuesdays in 2020 and 2021. I am very grateful that the Covid-period brought us together. In any type of conversation, you know the right things to say, whether it was about my research, society in general, or my personal life. I strongly believe you are the scholar that the economic science needs, but most of all you are an extremely good friend. Eda, I really enjoyed our time in room 509. Your kind soul, refreshing ideas and poetic attitude in life have lifted my spirits many times. I would also like to thank Maite, for unknowingly paving the way for me to start both the Research Master and subsequently the PhD. Mart, mainly for being witty and appreciative of my kind of humor, but also for many enlightening discussions we've had during the master and the early years of my PhD. Robin, for sharing many (virtual) coffees and experiences with similar research questions. And lastly, Melline, my PhD buddy from below the Rhine. I look back at some very nice academic trips together, talking about skills (we still need to write a paper together!), and about the great cultural divide between the north and south of the Netherlands.

Een proefschrift schrijven is meer dan teksten typen, het is vaak een traag, eenzaam en confronterend proces geweest. Daar heb ik de nodige ondersteuning bij gehad, veelal ook van buiten de wetenschap. Anna, ik bewonder je daadkracht, je positieve kijk op het leven en je aanstekelijke energie. Bedankt voor je luisterend oor en nuchtere perspectief - op nog vele jaren. Bertram, dank voor je vrije blik op de

wereld, die me stimuleerde om wat meer buiten de gebaande paden te gaan. Ik kijk met trots en tevredenheid terug op wat we met Hack Je Les hebben bereikt. Wineke, Gerrienke en Anneriek: dank voor hulp bij hoepeltjes, rustig opbouwen en maskers afzetten. Naomi, Ellen en de Zwolse voetbalmeiden: voor alle sportieve momenten, maar vooral voor al het andere eromheen. Kiki, how nice it was to work with you on the cover for this book – even though it has already been 8 years since our study abroad!

Lieve familie Timmerman, bedankt voor wie jullie zijn en hoe jullie me hebben opgenomen in jullie gezin. Albert, voor je oprechte interesse, het aanvoelen wanneer ik mezelf even niet kon motiveren en het laatste zetje bij het inleveren ("do it!"). Albertje, ik bewonder je liefde voor alles om je heen, hoe je kan genieten van kleine dingen ("allemaal gratis!") en hoe je tactvol kan corrigeren ("dan heb je zo'n mooie woordenschat Femke, en dan gebruik je zulke woorden..."). Jan en Aline, ik heb het gevoel er een extra broer en zus bij te hebben gekregen, daar ben ik heel dankbaar voor! Alex en Corine, ik kijk uit naar nog vele jaren delen hoe het is om met een Timmerman te gaan, ik hoop dat we met z'n allen nog een keer naar Bruce kunnen.

Lieve familie. Mijn lieve broer en zus, mijn mede-tijdreizigers, hoewel ik me waarschijnlijk mijn hele leven jullie grote zus zal voelen, zie ik ook steeds meer dat ik veel van jullie kan leren. Jelte, we lijken meer op elkaar dan we denken, en ik kijk ernaar uit om dat samen nog verder te ontdekken. Welmoed, ik zeg het niet altijd met veel woorden, maar ik koester de band die we hebben en heb bewondering voor hoe jij steeds meer je eigen pad durft te kiezen, zowel alleen als met Bob. Mama, van jou heb ik geleerd om van niets iets te maken: ik bewonder je creativiteit en je doorzettingsvermogen. Ik hoop er een stukje van in mezelf mee te dragen. Heit, ook al bewandel ik soms andere paden dan jij, ik voel altijd je steun, trots en bemoediging. Ik heb van je geleerd dat het leven één groot avontuur kan zijn, waarin je nieuwe dingen vooral moet proberen. Opa en oma, bedankt voor de vele open gesprekken over de betekenis van het leven, en al het mysterie daaromheen. Lieve beppe, ik mis uw stralende ogen en opgeheven hand. Bedankt voor alles.

En Marije. Voor wie je bent en wie je wordt. Duizendmaal dank voor je steun, je vertrouwen en je kritische vragen. Hoe donker het ook mag zijn, ik wens geen dag op deze wereld zonder jou om me heen. Ik kijk uit naar een gezond, oprecht en tevreden leven met je.

About the author

Femke Dina Cnossen was born on the 19th of September 1993, in Leeuwarden. She completed secondary school in Groningen, and obtained her bachelor degree in Economics and Business Economics at the Faculty of Economics and Business (FEB) at the University of Groningen. As part of the bachelor, she spent a semester studying abroad in Sendai, Japan. After graduating from the Research Master in Economics at the FEB in 2017, she continued to pursue a PhD at that same faculty, where she was supervised by prof. dr. Robert Inklaar, prof. dr. Steven Brakman and dr. Nikolova. This dissertation is the end result of that process.

During the Research Master, she participated in the 'Nationale Denktank' themed the Future of Vocational Education. During her PhD, she was co-owner of a spin-off project from the Denktank: Hack Je Les. The goal of this start-up was to improve the feedback culture at vocational schools throughout the Netherlands. She also was a coordinator of the Honours course Rebuilding Education in Groningen, which aimed at stimulating master students to develop ideas for the improvement of their own education. In early 2020, she spent a short research stay at the IZA in Bonn. Currently, she is working as a post-doctoral researcher at the department of Economic Geography, Faculty of Spatial Sciences in Groningen.