Economic Studies 222

Malin Backman

Essays on Women in the Labor Market

Department of Economics, Uppsala University

Visiting address: Kyrkogårdsgatan 10, Uppsala, Sweden
Postal address: Box 513, SE-751 20 Uppsala, Sweden
Telephone: +46 18 471 00 00
Telefax: +46 18 471 14 78
Internet: http://www.nek.uu.se/

ECONOMICS AT UPPSALA UNIVERSITY

The Department of Economics at Uppsala University has a long history. The first chair in Economics in the Nordic countries was instituted at Uppsala University in 1741.

The main focus of research at the department has varied over the years but has typically been oriented towards policy-relevant applied economics, including both theoretical and empirical studies. The currently most active areas of research can be grouped into six categories:

* Labour economics
* Public economics
* Macroeconomics
* Microeconometrics
* Environmental economics
* Housing and urban economics

Malin Backman

Essays on Women in the Labor Market

Technology, Inequality, and the Future of Work

UPPSALA
UNIVERSITET

Dissertation presented at Uppsala University to be publicly examined in Lecture Hall 2, Ekonomikum, Kyrkogårdsgatan 10, Friday, 7 February 2025 at 13:15 for the degree of Doctor of Philosophy. The examination will be conducted in English. Faculty examiner: Associate Professor Kristiina Huttunen (Aalto University, Department of Economics).

Abstract
Backman, M. 2024. Essays on Women in the Labor Market. Technology, Inequality, and the Future of Work. *Economic studies* 222. 182 pp. Uppsala: Department of Economics, Uppsala University. ISBN 978-91-506-3084-8.

Essay I: I study how digital skill requirements affect the gender composition of new hires in female-dominated administrative occupations in Sweden from 2006 to 2016. I track the prevalence of digital skill requirements in Swedish vacancy ads and analyze the gender composition of new hires through a matching process using Swedish register data. I also examine changes in age composition and its interaction with gender. Although men are overrepresented among computer-related educational fields and occupations, I find that the share of women has not decreased, despite a significant increase in digital skill requirements. Using employer fixed effects, I observe an increase in female hires ages 36-50 for ads that specify digital skill requirements. This may be due to the fact that increased exposure to technology increases the complexity of tasks, resulting in a higher demand for skilled workers, which are predominantly women in female-dominated occupations.

Essay II: I study the effects of robot and software technology exposure on the evolution of the gender wage gap in cities across the US from 1980 to 2010. Technology exposure, measured via task-replacing patents, has reduced labor demand in primarily male-dominated occupations. Consequently, the gender wage gap declined more in cities with high employment shares of exposed occupations. Moving from the 25th to the 75th percentile of city-level software exposure explains approximately 15% of the average decline in the gender wage gap in the sampled cities. The city-level association between robot patent exposure and changes in the gender wage gap is lower and less robust, possibly due to lower average wages for robot-exposed occupations. I perform a tentative analysis of AI-exposed occupations and show that AI patents are mostly related to tasks performed in male-dominated occupations high up in the wage distribution.

Essay III: (with Ola Andersson, Niklas Bengtsson, and Per Engström) Previous research suggests that student evaluations of teacher performance are biased against women. We test this hypothesis on economics students in a randomized, double-blind experiment, set up in a natural educational setting. During the Covid-19 pandemic, teaching assistants moved online and answered questions through email instead of on campus. We randomly assigned a male or female name to the instructions given by the online teachers. Importantly, the teachers actually responding to the questions did not know if they interacted with the students as male or female, which is a novel contribution to the literature. The course evaluation asked students to rate the mentors' helpfulness, knowledge, and response time. The results show no bias against the female mentor in any dimension. Our confidence interval around the zero effect does not overlap the effect sizes reported in highly influential previous studies.

Keywords: Technological Change, Gender Equality, Gender Wage Gap, Women's Labor Market, Student Evaluations

Malin Backman, Department of Economics, Box 513, Uppsala University, SE-75120 Uppsala, Sweden.

ISSN 0283-7668
ISBN 978-91-506-3084-8
URN urn:nbn:se:uu:diva-544385 (http://urn.kb.se/resolve?urn=urn:nbn:se:uu:diva-544385)

To Rut

Acknowledgements

I applied to do my bachelor's degree in Uppsala in early spring 2013. More than a decade later, I can confidently say it was one of the wisest decisions I have ever made. Over these years, I have completed my bachelor's, my master's, and now (hopefully) my Ph.D. at the Department of Economics in Uppsala. None of it would have been possible without the incredible talent, generosity, and kindness of the people with whom I have worked and studied here.

First and foremost, I would like to thank my main supervisor, Oskar Nordström Skans. Your dedication to your students is inspiring and I am very lucky to have been one of them. Without your thoroughness, wisdom, and support, it is unlikely that I would have had a thesis even half as good. I also would like to thank my assistant supervisor, Arizo Karimi. Your input has made all the difference in my work and you are the constant reminder that you truly can have it all, an enviable career in academia, and a fabulous life outside of work.

I want to thank my co-authors Per Engström, Niklas Bengtsson, and Ola Andersson. Working with you has been not only incredibly fun but also deeply informative. You have all been my teacher at some point before my PhD, and I wish I could tell my younger self that for a short stint, we were almost peers. I also want to extend my thanks to Julia from the administrative staff, whose invaluable support was essential for conducting our experiment in the first place.

For financial support, I want to acknowledge the project "The labor-market impact of firm-level adoption of AI and autonomous system" funded by the Marianne and Marcus Wallenberg Foundation. I also want to extend my thanks to the WASP-HS graduate school for the rewarding courses they supplied and the interesting discussion I was able to have with the other PhD students from different disciplines. I am also grateful for access to Swedish register data from IFAU, as well as public data from IPUMS and ILOSTAT.

I am very grateful to my discussants during my licenciate and final seminar, Adrian Adermon and Stefan Eriksson, for their meticulous reviews and invaluable feedback. Thanks also to members of the Uppsala Labor Group (ULG) and Uppsala Center for Labor Studies (UCLS) for all the helpful input I got along the way.

A heartfelt thanks goes to the real heroes of the Department of Economics, the administration team. Thank you for your professionalism, expertise, and, most importantly, your patience. I especially want to thank Ulrika and Nina, who have always had my back in any endeavor.

More important than any degree that Uppsala University may have granted me are the friends that I have gotten during my long tenure here. I want to thank my PhD cohort: Akib, Daniel, Erika, Gabriella, Hanfeng, Hiep, Lovisa, Qingyan, Rinni, and Zunyuan. I also thank the senior Ph.D. students who, beyond their friendship, offered invaluable advice and guidance. Thank you Anna, Josefin, Yaroslav, Elin, Anton, Sofia, Daniel, Edvin, Fei, Jan, Simon, Maximiliano, Zeynep, Raoul and Markus, to name but a few of everyone who has made the trip from Stockholm to Uppsala more than worthwhile. Venturing further back in time I want to thank Alva, Ellen, Emilie, and Kristin for holding my hand during my bachelor's. And thank you Anna, Anton, Elin, and Douglas for making my masters the unexpected extension of my teenage crimes.

I want to thank my parents, my in-laws, and my sister for their unconditional love and support. Lastly, I want to thank my husband and my baby. Thank you Alfons for being my best friend, my greatest cheerleader, and the unequivocal love of my life. And to Rut, out of all my accomplishments, you will always be my greatest.

Contents

Introduction

This thesis is written as part of the WASP-HS Graduate School that aims to broaden our understanding of how AI and autonomous systems affect humans and society at large. The graduate school and my PhD thesis were generously funded by the Marianne and Marcus Wallenberg Foundation. In my thesis, I study the different ways technological innovation has affected gender equality in the labor market, both historically from the 1980s and onwards and using more recent data. Understanding the motions through which novel technologies reshape the labor market for men and women, both combined and separately, allows us to make informed forecasts of the labor market of the future and the impact AI technology will have. In this Introduction, I summarize the motivation for these articles and place them in the broader context of technological change and the recent rise of artificial intelligence.

Throughout history, technological innovation has reshaped the labor market time and time again. Industrial robots significantly reduced the demand for labor in manufacturing. Similarly, personal computers drastically increased the productivity of clerical and analytical work. Most recently, the rise of artificial intelligence challenges our assumptions about which areas of the labor market remain uniquely human (Autor et al., 1998; Graetz and Michaels, 2018). Although distinct from prior labor-altering technologies, AI represents the latest chapter in the ongoing evolution of technological innovation and its impact on the labor market.

Gender-Specific Impacts of Technology

The impact of technological innovation spans virtually all parts of the labor markets in most countries. However, adoption rates during different eras differ across industries and occupations. Because of occupational gender segregation[1], men and women will therefore be differently exposed to these innovations. In addition, men and women may respond differently to the very same innovation. My thesis will study aspects related to both differences in exposures and responses.

The introduction of industrial robots in the late 1970s and early 1980s is an example of innovations that primarily affected male workers in the manufacturing industry, leading to declines in middle-wage workers with routine

[1]For more on occupational gender segregation, see for example Blau et al. (2013), Blackburn et al. (2000) and Hegewisch and Hartmann (2014).

jobs on factory floors and assembly lines that were easily replaced by robots (Blau and Kahn, 1997; Blau and Kahn, 2006). In contrast, the onset of information technology in the early 1900s primarily changed demand for female-dominated office and clerical jobs. Even earlier improvements in household production technology dramatically increased women's mobility outside of the home, with a direct impact on their ability to perform market work (Goldin, 2006).

The so-called task-based framework analyzes the impact on workers using a model where machines replace (or augment) tasks that are performed within certain occupations (Autor et al., 2003). For example, the task of assembling is one that has been replaced by industrial robots (Graetz and Michaels, 2018). This implies that the gender-specific consequences will depend on four things. First, what are the capabilities of the given technology, that is, what can it do and what tasks can it perform? Second, in which occupations are these tasks performed and in what intensity? Third, in which occupations are complementary tasks performed? Fourth, in which occupations do men and women work? Thus, to understand the implications of technology on gender equality and the gender wage gap, one must first establish what the technology can do, who is already doing that, and who will benefit from increased productivity in the tasks that the technology can perform.

Research shows that technological innovations have significantly influenced gender equality, a trend likely to continue with the rise of artificial intelligence. Most notably, part of the decline in the gender wage gap during the 1990s is attributed to the introduction of personal computers. This is due to its complementarity with female-dominated skills and tasks, such as social skills and analytical tasks (Spitz-Oener, 2006; Beaudry and Lewis, 2014; Cortes et al., 2018).

My thesis contributes to our understanding of these processes by i) analyzing the gender-specific responses to digital technologies on white collar work using Swedish data that match posted vacancies and register data and ii) study how the local exposure to technical innovations has affected the gender wage gap in US cities.

I next summarize these articles, and then return to a discussion regarding how AI-technologies may affect gender equality going forward. My thesis also contains a chapter on gender biases in student evaluations that I return to at the end of this introduction.

Essay I: Women still wanted? Digitalization and labor demand in female-dominated occupations: Evidence from Swedish Vacancy Ads

The first essay focuses on the increase in computer and / or software-related skill requirements in female-dominated clerical jobs and how it relates to the gender and age distribution among workers over time. To this end, I use Swedish register data that include individual-level information on gender, age, and occupation, among other things. I combine these data with the text of the posted vacancies from 2006 to 2016.

Because men are overrepresented in computer science-related educational fields, there is reason to believe that an increase in software reliance in certain occupations should increase the number of men in these roles (Lindley, 2012). I measure the digital skill requirements through simple text analysis of the vacancy ads and track the prevalence of certain software programs for a handful of occupations over time. I focus on female-dominated clerical jobs as they are prone to software program exposure but do not take part in the production of said technology. These occupations will also demonstrate whether changes in tasks related to technology adoption have the potential to make certain occupations less segregated.

For each position, I match the ad to a new hire in the register data based on occupation, location, employer, and month of entry. It is not possible to link a specific hired worker to a specific ad, bu I will interpret the content of the as reflecting the characteristics of the jobs performed in that specific establishment and occupation at that point in time. For each combination of establishment, occupation, and year, I calculate the share of women and the age distribution of new hires and relate it to the digital skill requirements in the posted ads for that exact combination of establishment, occupation, and year.

I focus the first part of my analysis on a comparison across jobs with the same occupation to isolate how men and women sort across employers who post different types of ads. I find that including digital skill requirements in the vacancy ad has no impact on the share of women among new hires. However, advertisements with digital skill requirements appear to increase the share of younger (≤ 35) workers.

To investigate whether the patterns are confounded by other differences across firms, I also estimate richer models that compare across occupations, within establishments. These models handle concerns that software-intensive jobs tend to be posted in firms that attract younger workers, and perhaps more men, in general. The results from these richer models indicate a positive relationship between ads with digital skill requirements and the share of women, within establishments. The effect, which is only marginally significant, is driven by prime-aged (thus, experienced) women.

Overall, none of my results indicates that the adoption of technology in these female-dominated clerical jobs caused a growing share of male workers, as would perhaps have been expected if male skills or preferences are complementary to technology. There are several potential mechanisms at work. One is that gender stereotyping is stronger than the changes in tasks. Changes in tasks may not be big enough to redefine who we think is most suitable for these types of clerical jobs, and instead, the women who inhabit them adapt and learn new skills. Another possible mechanism, related to the increase in women aged 36 to 50 years of age for digital positions, is that reliance on software technology increases the difficulty and abstractness of tasks (Barth et al., 2023). The positions are then more suitable for more experienced workers, which for these female-dominated occupations will be women, and slightly older.

The paper also points to interesting avenues for future research. The same method can be used in future research on other types of technology-related changes in labor demand. One such expansion could be to analyze the effects of labor-augmenting AI technology on the gender of new recruits. Are women and men equally likely to apply for and be recruited to roles that rely more heavily on AI tools?

Essay II: Software, Robots, AI & The Gender Wage Gap - Evidence from US Cities

The second essay focuses on robot and software exposure at the city level and how it relates to changes in the gender wage gap between 1980 and 2010 based on three distinct labor market phenomena. First of all, occupations differ greatly in their risk of automation due to robot and software technology innovation due to differences in the tasks that they perform. Workers on factory floors and assembly lines risk replacement by industrial robots, and certain types of operators and monitoring occupations risk replacement by software technology. Secondly, due to occupational gender segregation, men and women do not always work in the same occupations. This means that men and women are differently affected by the introduction of different types of technologies and face different risks of technology-induced labor displacement. Lastly, certain occupations are more common, that is, employ relatively more people in some cities compared to others. This means that technology-related changes in the overall structure of the labor market will differ between cities depending on the importance of certain occupations.

I use a patent-based exposure measurement created by Webb (2019) that assigns each occupation a percentile score based on the overlap between the occupations tasks as described in O*NET and the tasks performed by the patented technology. I find that robot exposure is high among primarily low-wage male workers and software exposure is high among primarily middle-

wage male workers. Occupations with high exposure rates declined in both employment shares and wages between 1980 and 2010.

In light of all this, I hypothesize that the gender wage gap declined more rapidly in cities with a high initial concentration of technology-exposed occupations, as these male-dominated roles experienced reductions in both employment and wages.

To test this hypothesis, I regress the change in the gender wage gap from 1980 to 2010 on the share of workers in the most (above median) exposed occupations in 1980. I find exactly what is hypothesized, the gender wage gap declined faster in cities with large shares of exposed workers in 1980, prior to the large scale introduction of both industrial robots and personal computers, compared to cities that were not as exposed. I attribute this to technology adoption decreasing the average relative wage for primarily male workers and/or displacing more male than female workers.

I do a tentative analysis regarding AI exposure, and although not enough time has passed to estimate changes in wages and employment, it is clear that AI exposure will again affect primarily male workers. However, it appears to affect mostly high-wage occupations instead of low- to middle-wage occupations that were most exposed to robot and software technology, respectively.

Labor Market Impact of AI Deployment

The deployment of AI is, from a labor perspective, still in an early phase in most parts. As a consequence, we are constrained to learning from other recent deployments of technology, and this is the approach taken in this thesis. These insights can hopefully teach us how to think about AI and the labor market of the future in a more structured and informed way.

What will happen to the labor market as AI technology becomes more commonplace? Again, we can take what we know about the capabilities of AI and look to the past to make some educated guesses. Previous technological revolutions revolved mainly around automating routine tasks, since they were easily programmable. Middle-wage occupations focused on routine tasks saw a decrease in labor share, while low-wage manual roles and high-wage analytical jobs increased in labor share (Goos et al., 2009; Goos et al., 2014).[2] AI technology has the ability to perform non-routine, abstract and analytical tasks that are mostly found in high-wage occupations that require college education(Brynjolfsson and Mitchell, 2017). Automating these tasks will decrease labor demand and likely compress the wage distribution. The net effect on the gender wage gap remains ambiguous and will depend on the type of AI technology that will be most popular. Generative AI appears to affect mostly

[2]Offshoring has also played a significant part in job polarization in Europe and the US, see Goos et al. (2014) for more on this.

female-dominated occupations but more general, yet simpler, AI, as used in Webb (2019), appears to affect mostly male-dominated occupations.

Beyond the displacing effects of AI technology, the expansion of AI tools will create jobs in the AI sector and related industries. Machine learning engineers and data scientists will be needed to continuously write and train models and integrate AI systems into different applications. The high volume of data needed will require more data engineers and database administrators to ensure access to high-quality data and that data regulations are followed. There will also be a need to facilitators that either improve the user-friendliness of AI applications or that train workers in using AI tools in their new augmented roles. All of these types of occupations and tasks will likely grow in size and are, for the most part, highly qualified roles. They are also mainly male-dominated and looking at the gender distribution of relevant STEM-related educational fields, they are likely to remain male-dominated. Not only will this mean that women might miss out on important changes in the labor market, women's needs and interests risk being overlooked when new AI models and tools are created.

On the other hand, AI tools will likely improve the flexibility of many roles regarding when and where they can be performed. This benefits mainly women who shoulder the bulk of child care and household duties (Goldin and Katz, 2016). This is not to be taken lightly, as Goldin (2014) refers to enhanced temporal flexibility as the last chapter of the 'grand gender convergence'.

Although the introduction of new technologies significantly alters labor market conditions, it is important to remember that workers and occupations are highly adaptable. Recent research has shown that occupational titles persist long after their task makeup has changed and technology-induced changes occur primarily within occupations (Atalay et al., 2020). Labor-augmenting technology, i.e., technology that improves worker productivity, increases labor demand even outside of technology-producing industries (Autor et al., 2024). However, this type of adaptability may require significant policy interventions, such as increased access to training.

One point to bear in mind when discussing the future is that occupational gender segregation influences who we see as most fit to inhabit certain occupations. Gender stereotyping runs deep and tends to move slower than technological innovation. For this reason, we might see smaller demographic changes than should be predicted considering the monumental influence of new technologies. I find that significant changes in technology-related skill requirements do not alter the share of women who perform a specific occupation.[3]

[3] See Essay I for more discussion on this topic.

Gender Differences in AI Exposure

A first step towards understanding the impact is to illustrate the gender-specific exposure to AI. AI technology has the potential to automate tasks in a broad array of occupations, but the gender-specific consequences will, as previously mentioned, depend on the gender composition of said occupations. It will also depend on which group of workers is most likely to implement AI tools and benefit from the increased productivity.

Webb (2019) quantifies AI exposure at the occupation level by assessing how tasks in the O*NET occupational dictionary overlap with tasks described in AI-related patents. O*NET details the task contents of almost 1000 occupations throughout the US economy and describes how common or important they are for any given occupation. Webb (2019) dissects each task into two-word phrases such as "analyze data" and "cook food" and calculates how common they are among patents for different types of technology. The more registered patents there are that perform an occupations task, the higher the exposure score, weighted by how important that task is for that occupation.[4] This measurement is considered to measure the extent to which AI can replace humans in performing a given occupation.

In figure 1, I relate the (predicted) AI-exposure of occupations to the share of women and the average wage in these occupations using US data from the American Community Survey (ACS) (Ruggles et al., 2021). When using the exposure measurement from Webb (2019), it appears that the risk of automation by AI technology is primarily high in male-dominated occupations and primarily high-wage occupations. This pattern is not unique to the US. In figure 2, I aggregate global employment data from 102 countries, adjusted for differences in female labor force participation, and find that AI exposure is still highest in occupations in which men tend to work.

[4] Similar approaches using patent data to calculate exposure to technology are found in Dechezleprêtre et al. (2021), Kogan et al. (2021), and Mann and Püttmann (2023) and most recently Autor et al. (2024).

Figure 1. AI Exposure - Female Share and Average Wage in the US

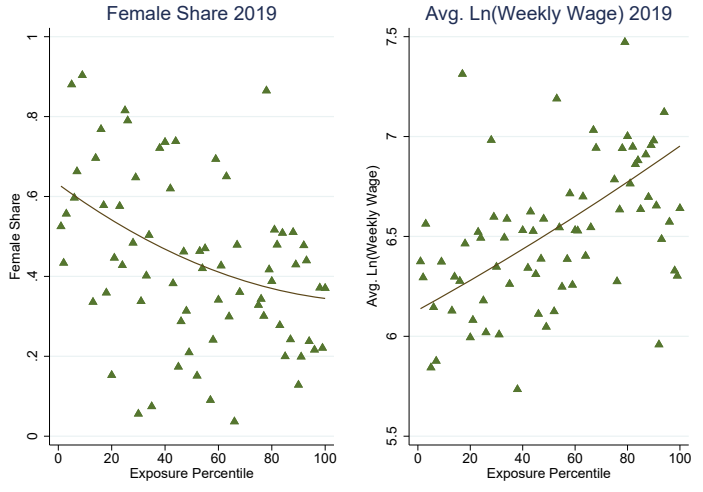

NOTE.— Left: AI Exposure and the female share of workers in occupations. Right: AI Exposure and the average Ln(Weekly Wage). Female share and average wage are calculated within each equally sized bin of exposure percentile, set at 100 bins. Data: Weighted data from American Community Survey (2019), restricted to employed individuals ages 16-64 living in a metropolitan area (AMA) (Ruggles et al., 2021). Wages are adjusted to full-time full-year equivalence, winsorized at the bottom 1% and adjusted to 2010 price levels. Exposure measurement from Webb (2019).

Figure 2. AI Exposure - Female Share using Global Aggregate

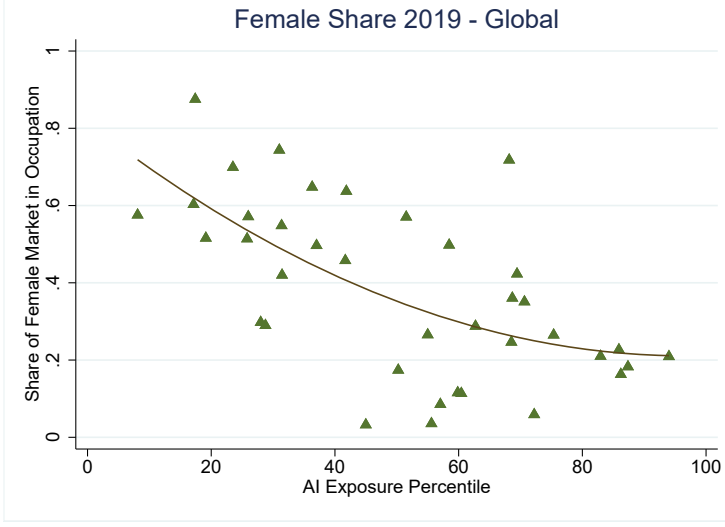

NOTE.— AI Exposure and the share of the female labor supply employed occupations using a global aggregate. Female share is calculated within each equally sized bin of exposure percentile, set at 100 bins. Data: Employment shares from ILOSTAT. Exposure measurement from Webb (2019).

However, the overall threat of AI exposure differs between countries. The left panel of figure 3 shows that the rate of exposure to AI differs noticeably between nations. The right panel of figure 3 shows that the same holds true for the male-female difference in AI exposure. Most notably, men are more exposed than women in all 102 countries in the data, except 5. In figure 4, I relate the country-level exposure rate and the male-female exposure difference to the Gross Domestic Product (GDP) per capita. It appears that there is a slight negative relationship between the overall AI exposure rate and GDP per capita. On the other hand, there is a clear positive relationship between the male-female exposure difference and GDP per capita. It is evident that AI will play different roles in terms of gender equality in different countries. In richer countries, with a higher GDP per capita, it appears that AI will have the most disparate effects on men and women.

Figure 3. Male-Female Difference AI Exposure

NOTE.— Left: Total weighted exposure scores for countries in the ILOSTAT database. The Graph includes the 10 countries with the highest and lowest relative exposure score, respectively. Red bar "TOT" displays the unweighted global average. Right: The weighted average AI exposure rate for male workers minus the equivalent for female workers for countries in the ILOSTAT database. The Graph includes the 10 countries with the largest and smallest male-female difference, respectively. Differences in female labor force participation is accounted for. Red bar "TOT" displays the unweighted global average. Data: Employment shares from 102 countries from ILOSTAT (2019). When 2019 is unavailable, data from 2018 or 2019 have been supplemented. Exposure measurement from Webb (2019).

Figure 4. Male-Female Difference AI Exposure

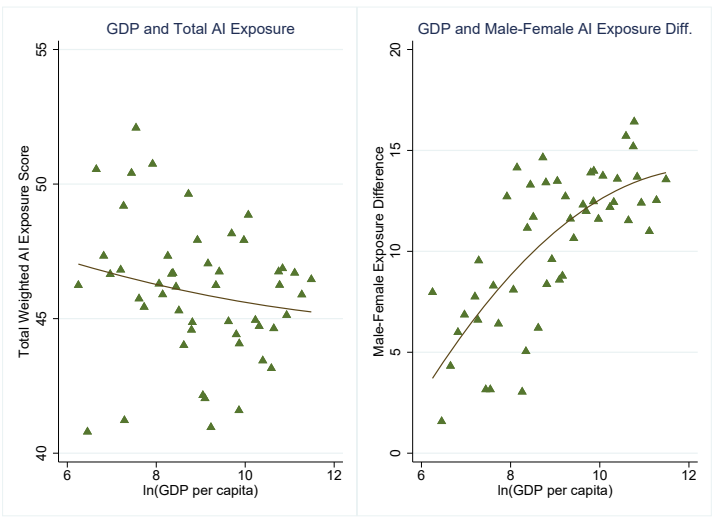

NOTE.— Left: Total weighted exposure scores over Ln(GDP). Total weighted exposure score is calculated within each equally sized bin set at 50 bins. Right: Male-Female AI Exposure difference over Ln(GDP). Average Male-Female Exposure difference is calculated within each equally sized bin, set at 50 bins. Data: Exposure rates, both total and male-female difference is calculated using employment shares from 102 countries from ILOSTAT (2019). When 2019 is unavailable, data from 2018 or 2019 have been supplemented. AI Exposure measurement from Webb (2019). Data on GDP from The World Bank, in current US Dollars.

It should be noted that the patents used for Webb (2019)'s exposure measurement only include those registered up to and including 2018 and will therefore not include today's more generative AI technology. It includes several types of technology that can perform analytical and planning tasks, as well as other types of algorithmic AI tools. For more recent AI technology, Felten et al. (2023) have created an exposure measurement based on generative AI technology. Felten et al. (2021) create this exposure by measuring the overlap between O*NET task descriptions and AI capabilities as described by the Electronic Frontier Foundation (EFF). In figure 5, I again relate the female share and average wage to the AI exposure rate of occupations, but this time I focus on language modeling and image-generating technology. In these instances, AI exposure is primarily high in female-dominated occupations, but the relationship is more clear with respect to language modeling. Similarly to the figures using the exposure measurement from Webb (2019), the exposure rates are primarily high in high-wage occupations.

From this we can gather that AI technology may not only affect other types of occupations than previous technological innovations, but different types of AI technologies may affect men and women differently. It is apparent that the wide range of tasks that AI technology is capable of performing, such as

Figure 5. Generative AI Exposure - Average Wage and Educational Background in the US

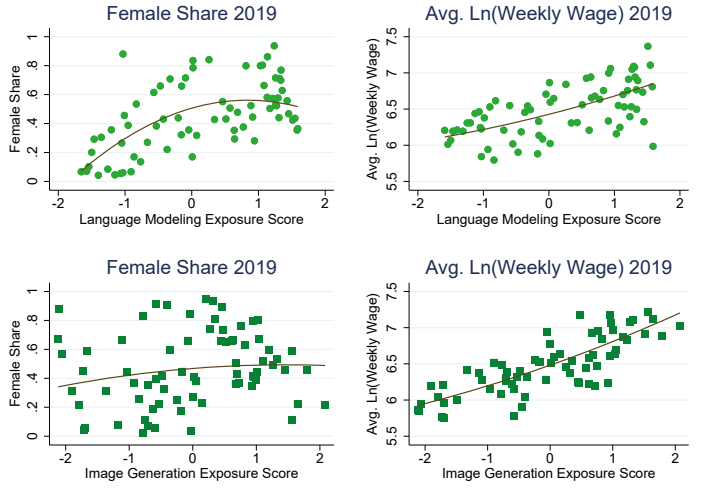

NOTE.— Top: Female share and average ln(weekly wage) over Language Modeling AI Exposure. Female share and average wage is calculated within each equally sized bin of exposure score, set at 100 bins. Bottom: Female share and average ln(weekly wage) over Image Generation AI Exposure. Female share and average wage is calculated within each equally sized bin of exposure score, set at 100 bins. Data: Weighted data from American Community Survey (2019), restricted to employed individuals ages 16-64 living in a metropolitan area (AMA) (Ruggles et al., 2021). Wages are adjusted to full-time full-year equivalence, winsorized at the bottom 1% and adjusted to 2010 price levels. Exposure measurement from Felten et al. (2023).

data analysis, content creation, and logistical planning, should be considered when discussing the overall results of AI deployment. We lose important heterogeneity if we always view AI as a technological monolith. A similar case can be made when we discuss the impact of computers on the labor market, as innovations in software technology broadened the capabilities of computers in the workplace (Atalay et al., 2018).

There is much uncertainty regarding the labor market of the future, but one thing is certain: AI technology will have played an important role in shaping it and will continue to play an important role in how it evolves beyond that. In light of this uncertainty, there is much to learn from researching previous technological revolutions. The specifics of the technology may change, but it is my belief that humans and firms remain largely the same. As we adapt and innovate, the lessons of the past remain our most valuable tool to navigate uncertainties of the future. I base my research on this principle and hope that my findings aid in making the future more clear.

Essay III: Are Economics Students Biased against Female Teachers? Evidence from a Randomized, Double-Blind Natural Field Experiment

The final essay emerged from my teaching tasks during the Ph.D. where we performed and experiment to study biases in student evaluations. Bias against female teachers in student evaluations has long been a major concern. Several studies for an extended period of time have repeatedly found that female teachers receive, on average, lower scores in student evaluations (Boring, 2017; Heffernan, 2021; Ceci et al., 2023). However, it is very difficult to correctly identify and measure the exact extent of any potential bias whilst simultaneously controlling for all aspects of teaching quality. This is when an experimental setting is needed.

We construct an experimental setting where we can vary the gender of the teacher whilst keeping the teaching quality constant. During the Covid-19 outbreak, Economics students at Uppsala University were encouraged to email in their questions in lieu of in-person study sessions and lectures. This was then formalized to perform the experiment during 2022 and students were asked to send all their questions to a designated email address. The questions were answered by me, but the signature was randomly chosen between a male-sounding name (Anton) and a female-sounding name (Elin). The students did not know that all the answers came from the same source and, importantly, I did not know what the signature would be when I wrote the answer. At the end of the semester, the students were asked to evaluate their "email teachers" on three dimensions: 'helpful,' 'knowledgeable,' and 'timely.'

We find no significant differences in the average scores for the male and female teachers, on any dimension or when we average the three. In particular, our estimates are very small and precise, which differs noticeably from previous research. We attribute this to the experimental nature of our study that allows us to control for teaching quality, especially nuances in teaching style that can affect evaluations but are hard to directly observe. Secondly, our results differ from other experimental papers but ours is the first, to our knowledge, to employ a double-blind strategy in which gender assignment is both random and unknown to all participants.[5]

We acknowledge that online teaching of this nature is very different from in-person teaching and that our experiment, had it been possible, would have benefited from increased saliency in gender exposure. With new generative AI, it may be possible to perform a similar experiment with AI-generated online teachers in the near future.

[5] See MacNell et al. (2015) for the study most similar to ours, but not double-blind.

References

Atalay, Enghin, Phai Phongthiengtham, Sebastian Sotelo, and Daniel Tannenbaum (2018). "New technologies and the labor market". *Journal of Monetary Economics* 97, 48–67.

Atalay, Enghin, Phai Phongthiengtham, Sebastian Sotelo, and Daniel Tannenbaum (2020). "The evolution of work in the United States". *American Economic Journal: Applied Economics* 12.2, 1–34.

Autor, David, Caroline Chin, Anna Salomons, and Bryan Seegmiller (2024). "New frontiers: The origins and content of new work, 1940–2018". *The Quarterly Journal of Economics*, qjae008.

Autor, David H, Lawrence F Katz, and Alan B Krueger (1998). "Computing inequality: Have computers changed the labor market?" *The Quarterly journal of economics* 113.4, 1169–1213.

Autor, David H, Frank Levy, and Richard J Murnane (2003). "The skill content of recent technological change: An empirical exploration". *The Quarterly journal of economics* 118.4, 1279–1333.

Barth, Erling, James C Davis, Richard B Freeman, and Kristina McElheran (2023). "Twisting the demand curve: Digitalization and the older workforce". *Journal of Econometrics* 233.2, 443–467.

Beaudry, Paul and Ethan Lewis (2014). "Do male-female wage differentials reflect differences in the return to skill? Cross-city evidence from 1980-2000". *American Economic Journal: Applied Economics* 6.2, 178–94.

Blackburn, Robert M, Jennifer Jarman, and Bradley Brooks (2000). "The puzzle of gender segregation and inequality: A cross-national analysis". *European Sociological Review* 16.2, 119–135.

Blau, Francine D, Peter Brummund, and Albert Yung-Hsu Liu (2013). "Trends in occupational segregation by gender 1970–2009: Adjusting for the impact of changes in the occupational coding system". *Demography* 50.2, 471–492.

Blau, Francine D and Lawrence M Kahn (1997). "Swimming upstream: Trends in the gender wage differential in the 1980s". *Journal of labor Economics* 15.1, Part 1, 1–42.

Blau, Francine D and Lawrence M Kahn (2006). "The US gender pay gap in the 1990s: Slowing convergence". *Ilr Review* 60.1, 45–66.

Boring, Anne (2017). "Gender biases in student evaluations of teaching". *Journal of public economics* 145, 27–41.

Brynjolfsson, Erik and Tom Mitchell (2017). "What can machine learning do? Workforce implications". *Science* 358.6370, 1530–1534.

Ceci, Stephen J, Shulamit Kahn, and Wendy M Williams (2023). “Exploring gender bias in six key domains of academic science: An adversarial collaboration”. *Psychological Science in the Public Interest* 24.1, 15–73.

Cortes, Guido Matias, Nir Jaimovich, and Henry E Siu (2018). “The "end of men" and rise of women in the high-skilled labor market”. Tech. rep. National Bureau of Economic Research.

Dechezleprêtre, Antoine, David Hémous, Morten Olsen, and Carlo Zanella (2021). “Induced automation: Evidence from firm-level patent data”. *University of Zurich, Department of Economics, Working Paper* 384.

Felten, Edward, Manav Raj, and Robert Seamans (2021). “Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses”. *Strategic Management Journal* 42.12, 2195–2217.

Felten, Edward W, Manav Raj, and Robert Seamans (2023). “Occupational heterogeneity in exposure to generative AI”. *Available at SSRN 4414065*.

Goldin, Claudia (2006). “The quiet revolution that transformed women’s employment, education, and family”. *American economic review* 96.2, 1–21.

Goldin, Claudia (2014). “A grand gender convergence: Its last chapter”. *American economic review* 104.4, 1091–1119.

Goldin, Claudia and Lawrence F Katz (2016). “A most egalitarian profession: Pharmacy and the evolution of a family-friendly occupation”. *Journal of Labor Economics* 34.3, 705–746.

Goos, Maarten, Alan Manning, and Anna Salomons (2009). “Job polarization in Europe”. *American economic review* 99.2, 58–63.

Goos, Maarten, Alan Manning, and Anna Salomons (2014). “Explaining job polarization: Routine-biased technological change and offshoring”. *American economic review* 104.8, 2509–2526.

Graetz, Georg and Guy Michaels (2018). “Robots at work”. *Review of Economics and Statistics* 100.5, 753–768.

Heffernan, Troy (2021). “Sexism, racism, prejudice, and bias: A literature review and synthesis of research surrounding student evaluations of courses and teaching”. *Assessment & Evaluation in Higher Education* 47.1, 144–154.

Hegewisch, Ariane and Heidi Hartmann (2014). “Occupational segregation and the gender wage gap: A job half done”. Tech. rep. Institute for Women’s Policy Research.

Kogan, Leonid, Dimitris Papanikolaou, Lawrence DW Schmidt, and Bryan Seegmiller (2021). *Technology-skill complementarity and labor displacement: Evidence from linking two centuries of patents with occupations*. National Bureau of Economic Research.

Lindley, Joanne (2012). “The gender dimension of technical change and the role of task inputs”. *Labour Economics* 19.4, 516–526.

MacNell, Lillian, Adam Driscoll, and Andrea N Hunt (2015). “What’s in a name: Exposing gender bias in student ratings of teaching”. *Innovative Higher Education* 40.4, 291–303.

Mann, Katja and Lukas Püttmann (2023). “Benign effects of automation: New evidence from patent texts”. *Review of Economics and Statistics* 105.3, 562–579.

Ruggles, Steven, Sarah Flood, Sophia Foster, Ronald Goeken, Megan Schouweiler Jose Pacas, and Matthew Sobek (2021). “PUMS USA: Version 11.0 [dataset]”.

Spitz-Oener, Alexandra (2006). “Technical change, job tasks, and rising educational demands: Looking outside the wage structure”. *Journal of labor economics* 24.2, 235–270.

Webb, Michael (2019). “The impact of artificial intelligence on the labor market”. *Available at SSRN 3482150*.

Essay I. Digitalisation and Labor Demand in Female-Dominated Occupations: Evidence from Swedish Vacancy Ads

Acknowledgements: I would like to express my gratitude to Oskar Nordström Skans, Arizo Karimi, as well as the seminar participants at the Uppsala Center for Labor Studies and the Labor Group at Uppsala University for their helpful comments and suggestions. This work was funded by the Wallenberg AI, Autonomous Systems and Software Program – Humanities and Society (WASP-HS) funded by the Marianne and Marcus Wallenberg Foundation.

1 Introduction

With recent technological advancements, such as increasing digitalisation and software-related advancements, large shifts have occurred in the demand for certain types of skills and, by extension, certain types of workers (Autor et al., 2003; Atalay et al., 2018; Spitz-Oener, 2006). Despite progress, the labor market in most countries remains distinctly gender segregated across jobs and occupations, a process that begins with the choice of educational fields. Fields like engineering and computer programming remain predominantly male-dominated (Lindley, 2012). In a labor market dependent on digital technology, these are highly sought-after skills. As occupations require more experience with software-related technology, this could increase the demand for male workers. However, at the same time, recent technological developments and the automation of tasks also appear to increase the relative demand for interpersonal and communicative skills and stereotypes, as well as some recent research, suggest that women, more than men, excel at these types of skills (Cortes et al., 2018). This study aims to analyze whether male workers replace female workers in female-dominated administrative occupations as technology-related competencies increase. I approach this question by analyzing whether or not more male candidates are hired when vacancy ads include digital skill requirements. I also look at the interaction between age and gender, to more comprehensively capture changes in labor demand as a response to digital technology adoption.

To this end, I leverage the rich text data of Swedish vacancy ad posting from 2006 to 2016, matched to firm-level data on new hires, focusing on administrative, female-dominated, occupations. I focus on female-dominated administrative occupations because they have high levels of variation in skill requirements and task composition and are likely to be affected by the emergence of easily identifiable software and computer-related technologies (Deming and Kahn, 2018). When the level of software-oriented skill requirements increase, what happens to the demographics of these occupations? By applying simple text analysis to the vacancy ads, I extract the level of digitalisation by tracing the occurrence of technology-related words and phrases, stemming either from descriptions of the occupational role or from stated skill requirements. Combined with Swedish administrative data, the goal is to trace the digitalisation rates of vacant positions and the gender composition of new recruits at the firm-occupation level, also taking the age of the candidates into consideration. I am unable to separate supply from demand effects, and the hiring decision and recruitment are ultimately the net of these two effects.

The onset of computerisation in the labor market in the early 1980s has been shown to drastically change the task contents and skill requirements for certain occupations (Spitz-Oener, 2006). However, the specific contents and capabilities of personal computers have been somewhat overlooked and instead seen as a technological monolith. Atalay et al. (2018) show that the specific type of

software or ICT does matter and that different types of software have different effects on task contents and skill prices. For example, although the introduction of ICTs broadly increases the level of non-routine tasks in an occupation, the Microsoft Office Suite was shown to have the opposite effect. The rich text data of vacancy ads allow for a more in-depth analysis that takes the specifics of software technology into account. Parallel to the increased computerisation of workplaces, the rapid decline in the gender wage gap from 1980s and onward has, in part, been attributed to the complementarity of female skills to personal computers. Specifically, personal computers are seen as complementary to cognitive skills, common among female workers, and have been shown to increase the analytical skill contents of women's work (Black and Spitz-Oener, 2010; Beaudry and Lewis, 2014). The increased advancements of software technology, beyond the introduction of personal computers, should be seen as the next phase of the computerisation of workplaces and could mean different outcomes for the female labor supply, compared to previous phases.

My project relates to an international literature that use the contents of vacancy ads to inquire about skill requirements for certain occupations and how they relate to aggregate labor outcomes, such as skill prices or task composition (Deming and Kahn (2018) and Acemoglu et al. (2020)). However, the strength of my data is that I do not need to rely on aggregate labor outcomes and can instead study firm- or workplace-level recruitment and how it changes over time, both within firm and within occupation. To the best of my knowledge, I am the first to relate the stated software-related skill requirements to the gender composition among new recruits, using direct matching between vacancy ads and hiring firms. My work also contributes to the literature that relates labor market changes to changes in skill requirements in vacancy ads, showing that technological advancements create demand for new types of skills (Atalay et al., 2020). I am able to contribute to this literature by studying the demographic consequences of these demand changes, both at the aggregate and firm level.

The results show that, although the occurrence of digital skill requirements has significantly increased for several of the occupations in my working data, the share of women among new hires remains unaffected and the share of young workers (aged 35 or under) has increased. Including establishment fixed effects controls for both the tendency to adopt and implement technology and the tendency to recruit either women or younger people, for example. For these specifications, the results show that the share of women among new hires increases with the inclusion of digital skill requirements, specifically women of prime working age of 36 to 50 years of age. Overall, the results show that the digitalisation of occupations and tasks does not significantly shift the gender composition within occupations. Instead, it appears that increased digital skill requirements increase demand for experienced workers, which, on average, means women in these female-dominated occupations.

The remainder of the paper is organized as follows. Section 2 presents a general theoretical framework for how digital skill requirements affect the characteristics of new recruits. Section 3 describes the data sources, the measurements created, and the matching process of vacancy ads to new recruits. Section 4 describes the empirical approach, and Section 5 presents the results. The paper ends with a discussion of the results and conclusions in Section 6 followed by the Appendix.

2 Technological Advancements, Vacancy Ads and Recruitment

In this section, I describe how technological advancements affect the skill requirements that enter into vacancy ads and how this may affect the characteristics of new recruits, both in terms of gender, as well as age/experience.

When firms/workplaces adopt new digital technology, as they become available and/or financially viable, this shifts the task contents and skill requirements for certain occupations that the firm employs (Autor et al., 2003; Autor, 2013). These shifts enter into the vacancy ads, through changes in the overall task composition and/or changes in the stated digital skill requirements. These requirements often specify certain software expertise, programming proficiency, or data-oriented education. This paper focuses on technology and software-related skill requirements, but the overall changes in task content or skill requirements associated with digital technology adoption, such as certain tasks being automated and no longer included in the job description, are seen as part of the total effect even when not directly measured.

The inclusion of digital skill requirements may affect the pool of workers who apply for a specific opening. It may also affect the "ideal candidate" in the eyes of the recruiting firm and the criteria for whom they would most like to hire. These two processes combined generate specific characteristics of the new hire, such as gender. I am able to connect the contents of the vacancy ads to the characteristics of new hires, but I am not able to separate changes in the applicant pool from changes in the preferences of the recruiting firm.

Both of these processes have potentially gendered components. For example, requiring a specific educational background affects the gender composition of eligible candidates through educational gender segregation (Lindley, 2012). The demand-side effects on gender may also work through more subtle channels, where the increase in digital skill requirements is not enough to require a specific educational background, but enough to affect the general view of who is most suited for the position. Stereotypes predict that men are more able to handle more advanced software technology should the position require it. Stereotypes also predict women being more fit for positions that require more social and interpersonal skills (Cortes et al., 2018).

From the labor supply side, increasing the skills and experience demanded in vacancy ads evokes the issue of "overconfidence of men", which predicts that men are more likely to apply to ads than their equally experienced female counterparts (Bertrand, 2011). On the other hand, improvements to the digital infrastructure of firms mediate the possibility of more flexible working conditions that will benefit women, more so than men, who shoulder the greater share of home life responsibilities (Goldin, 2014). This could potentially increase the share of women who apply for such positions.

In addition to the gendered components of the application and recruitment processes, there is also reason to believe that the age of the candidate will matter. Younger workers are more exposed to new software from their more recent studies and/or from having grown up in a more digital age, which means that they could be more desired as firms introduce new software.[1] This would shift labor demand towards the younger applicant pool. In addition, recent studies have shown that firms that deploy software technology are less likely to hire older workers because of the cost of training and shorter employment time (Schleife, 2006). Van Borm et al. (2021) find that the older age of a job applicant signals to recruiters lower technology skills and trainability. On the other hand, since recent generations of software have improved in regards to the simplicity of their interface, it has lowered the threshold to run and interpret output. This would benefit more experienced workers who now have more improved tools to complement their knowledge (Barth et al., 2020; Zolas et al., 2021). Increased digitalisation has also been shown to increase the share of non-routine analytical tasks within occupations, and task complexity has been shown to increase with computer use (Atalay et al., 2018; Spitz-Oener, 2006). This could also shift demand towards more experienced workers, as the new, higher, level of abstraction requires more skilled workers. Due to high levels of occupational segregation, the desire for more experienced workers will implicitly mean the desire for specifically male or female workers, depending on the occupation.

3 Data and Measurements

In this section, I describe the data that I use to track both the digitalisation rate of occupations and the characteristics of new hires. I also describe how I define digitalisation in the vacancy ad texts and how I identify a new hire in the register data.

Vacancy ads

The vacancy ad data includes all posted ads from Arbetsförmedlingen (The Swedish Public Employment Service) from 2006 to 2016 and include ads for

[1]The average age at graduation is 28 in Sweden as of 2021. For more information, see https://www.ekonomifakta.se/fakta/utbildning-och-forskning/utbildningsniva/examensalder/

a wide variety of occupations and industries. As of now, only data up until and including 2016 are available. The number of ads vary between years and increases over time from around 200 000 to 700 000 ads per year. Each ad entry contains information about the employer, occupation, location of position (municipality), date the ad was posted, last day of application and the full ad text as it was posted. The ad text is used to determine the level of digital skill requirements and the identifying factors are used to match the ad with a potential new hire.

Using vacancy ads generally, and vacancy ads from Arbetsförmedlingen specifically, does create some externality issues. Firstly, the ads overrepresent occupations, firms, and industries that have high turnover rates and / or are expanding because they require new hires more frequently (Kuhn and Shen, 2013). Secondly, the ads from Arbetsförmedlingen will disproportionately represent the public sector, as they are required to announce vacancies as broadly as possible (Jensen, 2020). Table A1 shows that among the top 40 most common occupations, several of them overlap between the vacancy data and the employee data. However, the relative occurrence of certain occupations differ significantly. Table A2 shows similar tendencies for the industry distribution.

On the other hand, there are significant gains to using vacancy ads as opposed to occupational databases such as O*NET or DOT. Firstly, vacancy ads are updated more frequently than these databases and capture small changes that occur frequently throughout the studied time period (Jensen, 2020). Secondly, since the occupational databases describe what occupations *generally* do, it will miss within-occupation variation in skill requirements and task composition. Deming and Kahn (2018) find that skill requirements vary more between firms than between occupations, an important source of heterogeneity that is otherwise lost.

Register data

When combined, the register data documents the employer, occupation, place of work, municipality, and month of employment for each worker, as well as gender, age, and educational background. The register data includes all workers in the public sector and a sample of workers in the private sector. For all workers in the register data, around 50% have a missing occupational code due to sampling. However, the overrepresentation of public sector employers and large firms in the vacancy ad data means that the sampling rate remains higher for my working data. When limiting the employers to those found in my working data, only 35% have a missing occupational code.

Working data

In this article, I focus on the prevalence of digital technology and limit my scope to occupations that are most likely to be affected by this technology in a measurable way. Firstly, I focus on occupations that use digital technologies

as tools, rather than direct replacements, such that they show up in the vacancy ads when the skill requirements are described. Vacancy ad texts will not be able to capture the rate of technology exposure in the event that individual tasks or entire occupations are replaced by digital technologies. For this reason, I focus on administrative and clerical white-collar occupations. Related to this, white collar occupations are likely to have a broader scope for important variation in skill and technology requirements, compared to blue collar occupations (Deming and Kahn, 2018). Secondly, I want to avoid occupations that in one way or another are suspected to be creators of digital technology, which is why I exclude engineers of all kinds. Lastly, in the spirit of Goldin and Katz (2016), I focus on female-dominated occupations to more closely follow labor outcomes for women in the digital age.

The working sample is limited to the 12 most common, in terms of number of ads, occupations that satisfy these requirements. [2] The occupations used in the final sample are as follows:

- Controller
- Human Resources
- PR Consultant
- Administrator, Public Sector
- Social Secretary
- Administrative Assistant
- Accountant
- Aid Officer, Public Sector
- Office Secretary
- Financial Assistant
- Office Assistant
- Receptionist

Defining digitalisation

The digitalisation rate of a specific position is based on the content of the vacancy advertisements and mentions of technological skill requirements. This involves parsing ad text and marking it as digital if it mentions any of the predetermined technologies. The text analysis requires exact matching with respect to spelling but not in regards to capitalisation or punctuation. Deploying a fuzzy match is not considered necessary, since the exact matching appears to capture a consistent level of digitalisation and does not introduce the risk of faulty identification.

The list of technologies used is collected from Job Tech Development, a project lead by Arbetsförmedlingen (Swedish Public Employment Service), that aims to build a common infrastructure for digital job application and matching. These technology lists are based on interviews with industry organizations and / or gathered from the ad contents themselves. To verify the usefulness of these lists in this context, I manually read 500 ads for accountants and recorded any mention of digital and/or technological skill requirements that I could recognize. The resulting list of technologies overlapped significantly

[2]Table A1 shows the 40 most common occupations in the ad data, aggregated over all years from 2006 to 2016. See Table A10 for the female share of each occupation, before and after matching with vacancy ads.

with that of the corresponding Job Tech Development list of technologies commonly used by accountants. The main difference between the lists is that the manually created one includes common misspellings and Swedish compound words ('excel skills' = 'excelkunskaper'). Another significant difference between the lists is that the JobTech Development list is more comprehensive and includes less common technologies that were not included in the ads that I read and includes technologies that I failed to recognize as such. The Job Tech Development lists are chosen because of this and because they are available for any and all occupations included in the vacancy data, such that the analysis can be expanded to include more occupations for future analysis.

Some adjustments are made to the lists to accommodate the language most commonly used in vacancy ads. For example, the accounting software SAP R/3 is often referred to as either just SAP or R/3, and therefore all three iterations of this are included in the search list. Some digital mentions, such as the network classification MAN (Metropolitan Area Network), are removed from the list because it is a homonym with the very common Swedish word for 'man' and for 'one'.

Previous work using vacancy ads often uses a predetermined list of characteristics when parsing the text, most commonly a categorized set of skills/tasks to investigate the extent of routine-biased technological change (Deming and Kahn, 2018; Jensen, 2020; Spitz-Oener, 2006). Adding to this, using a predetermined and exhaustive list of technologies, defined clinically, means that the rate of digitalisation as it is measured is not sensitive to trend changes in wording or technology use outside the scope of this paper. See Table A4 for the full list of digital technologies considered in this paper. The lists of labor technologies are similar to those used in Hershbein and Kahn (2018).

If one or more of the labor technologies is found in the ad text, the ad is considered digital at the indicator level (similar to Jensen (2020)). I repeat this exercise whilst leaving out the most common technology for each occupation and year, to identify vacancy ads that stand out in terms of their digital skill requirements.[3] Figure 1 shows the digitalisation rate for each occupation over the sampled years for both definitions.[4] The figure shows that the level of digitalisation varies significantly between the occupations in the working data. Most occupations have experienced an increase in the digital skill requirements, whilst some have remained stable throughout the study period. The types of technologies that have increased are primarily the Microsoft Office Suite, accounting software (Fortnox, Hogia AB), resource planning software (Agresso, SAP R/3), and some database management software (SQL, Microsoft SharePoint). I do not find significant increases in more advanced or open-ended skill requirements such as Python or C++ programming, specific digital creation software, or experience with certain operating systems.

[3] See table A11 for the most common technology for each occupation and year.

[4] For the digitalisation rate of all occupations, see table A5.

Figure 1. Digitalisation Rate of Occupations

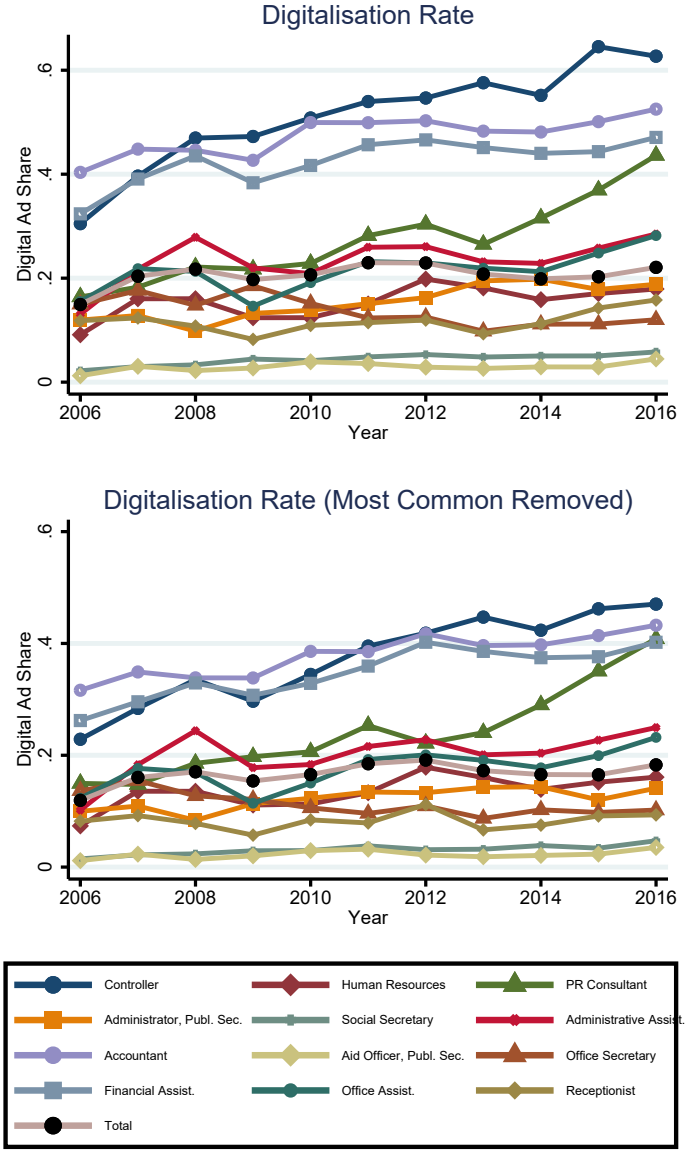

NOTE.—Digitalisation rate over occupations from 2006-2016, prior to matching with new hires.

I use a binary definition simply for ease of interpretation, but I include a continuous measurement, based on the number of digital skill requirements mentioned in the ad text, among the robustness tests.[5]

Previous work has shown that the task composition / skill requirements of occupations change along with the implementation of technology (Spitz-Oener, 2006; Atalay et al., 2020). Analyzing the specifics of these changes is beyond the scope of this paper and is instead considered part of the total effect.

Matching Vacancy Ads and New Hires

When each ad is classified as digital or non-digital, each ad is then matched to a potential new hire based on employer, county, occupation, and time of hire. The time of hire is denoted in months and a new hire is defined in one of two ways. Firstly, anyone with a starting month other than January is considered a new hire. This is because incumbent employees always have January as their starting month. Secondly, I must distinguish which employees with January as their starting month is a new hire. For this, I consider anyone who did not have an employment during the previous year or those who have switched employer from the previous year. However, I do discard any switchers that come from an employer that at least 30% of their current employees also come from, as these new hires are most likely the result of mergers with an already existing firm. The four identifying factors, employer, county, occupation, and time of hire, are assumed enough to, in the best case, match each new hire with the ad that they responded to and, in the second best case, find the new hires that were hired during the same time period that the firm had specific digital skill requirements for that type of position.

One component in the matching of ads to potential new hires that can be altered is the time window for which one could consider that the new hire is associated with that particular ad. The ad data contain two time stamps, the date of publication and the last date of application. The first cut-off point is rudimentary; anyone hired before the date of publication is not considered to be associated with the ad. You could also consider removing all new hires that were hired before the last date of application. However, positions are often open for extended periods, but may also be filled prematurely. The second cut-off point requires more consideration. The average notice period, the stipulated time between the notification of dismissal/termination and an employee's last day of work, for both employers and employees in Sweden is between 1 and 3 months and is governed by the employees collective agreement. If such an agreement is not in place, Swedish labor laws require a minimum of a 1 month notice period. Together, this means that the window should be at least 3 months from the last date of application. Testing the validity of this, I focus on the one-to-one matches in my data, contingent on being hired af-

[5]For the average number of digital skill requirements for each occupation and year, see figure A1 and table A12.

ter the ad was published. Table A6 shows that setting the hiring window to 6 months captures 90% of all matches, regardless of the matching criteria.

Table 1 shows the share of ads and new hires that are matched to at least one potential candidate for each group of matching criteria and for a three- and six-month hiring window. Matching on employer, occupation, and municipality is the closest the data allows to match on occupation and workplace, since workplace is not stated in the vacancy ad data. However, there is a significant discrepancy between the two data sources with respect to the location of the job. To remedy this, I opt for a more coarse geographical identifier, county. As discussed above, a hiring window of 6 months seemingly captures the majority of viable matches. Matching on the employer x county x occupation level, using a six-month hiring window, allows me to match 40% of all ads and 23% of new hires. Compared to previous studies using matched vacancy ads and new hires, this is a high match rate. Deming and Kahn (2018) are able to match 30% of ads and Jensen (2020) is able to match around 50% of all ads and slightly below 30% of new hires using employer*occupation as the identifying unit, ignoring any geographical matching. Several robustness tests are performed, altering either the matching granularity or the hiring window, to make sure that the conclusions are not driven by the matching criteria.

Table 1. *Matching Rate Summary - Vacancy and New Hires*

	Matching Granularity		
Hiring Window	Employer x Occupation	Employer x County x Occupation	Employer x Municipality x Occupation x
6 months	Ads: 45% New Hires: 29%	Ads: 40% New Hires: 23%	Ads: 38% New Hires: 20%
3 months	Ads: 42% New Hires: 27%	Ads: 38% New Hires: 22%	Ads: 35% New Hires: 19%

NOTE.— Matching rates for each granular level of matching, aggregated over all years 2006-2016.

New hires that are not matched to a vacancy ad are most likely hired through another channel, such as internal recruitment, recommendations, connections, or other ad posting site. The success rate of vacancies is not available in the data and it is not possible to know if a person was hired through Arbetsförmedlingen or some other channel. Between 10 and 15 % of the newly hired workers report that they have found their new position through Arbetsförmedlingen (SCB, 2017). Kettemann et al. (2018) shows, when using Austrian vacancy ad data, that around 14% of all advertisements fail to produce a new hire.

Table A7 shows that the match rates differ significantly between occupations. Some occupations may have longer recruitment times whilst others

fill their positions more quickly. It could be the case that the window is too small for some occupations. The matching rate is also lower when I include a geographical indicator, either a county or a municipality. For some occupations, remote work is more common, such that the location of the hiring firm/workplace does not match the location as stated in the ad. It could also be the case that the stated location of the firm in the register data simply does not match the location for where the position is intended to be placed. The matching process generates several duplicate matches, where a vacancy ad is matched to more than one potential candidate based on the matching criteria.[6] Duplicate matches can occur for several reasons. Some positions with the same occupation code at the same firm may be hired simultaneously but through different ads posted at different places. All such hires will therefore be matches to the ad that is available to me in Arbetsförmedlingens data.

A concern in regards to the matching is that the failure to match ads to new hires is skewed in terms of digitalisation and/or the share of women among new hires. To assess the degree of skewness, I compare both the share of women among new hires and the share of digital ads before and after matching. Table A9 shows that the digitalisation rate is lower among the matched sample compared to the unmatched sample. Table A10 shows that the share of women and the average age among new hires remain fairly similar in the matched and unmatched sample. As a robustness test, I match based on employer, county, and occupation rather than individual vacancy ads and new hires and redo the main analysis. This addresses, for example, duplicate matches and missed matches due to wrongly specified hiring windows. However, it does weaken the relationship between the content of the ads and the characteristics of the new hires as it will contain measurements from non-matched employees.

4 Empirical Model

To estimate the effect of digital skill requirements on the share of women among new hires and the age composition, I consider two variations of a simple model. The first considers within-occupation variation in the share of ads with digital skill requirements, and the second considers within-employer variation. The former will reflect how the demographics among new hires in these administrative occupations change when digital skill requirements enter vacancy ad texts. However, it will not capture how specific employers or workplaces adapt their recruiting when including digital skill requirements for their vacant positions. The latter model, which includes employer fixed effects, will capture how digital skill requirements affect the characteristics for people hired for specific positions, and not just occupations in general. Employers who extensively include digital skill requirements in their vacancy ads may

[6] See table A8 for the average number of duplicates per ad for each occupation

differ from other employers, in particular, in their propensity to hire women and workers of a certain age.

For each employer-county-occupation interaction, I take the yearly mean of all dependent and independent variables. I do this primarily to smooth out variation in digital skill requirements and employee characteristics since multiple hires are matched to multiple vacancy ads and vice versa. Yearly means at the employer-county-occupation level smooth out some of the skewness in either digitalisation rates or demographics among new hires created by these duplicate matches. It is not vital that each ad is matched to the exact new hire that responded to that specific ad but rather that I match them as closely as possible. This is based on the assumption that the inclusion of increased digital requirements in an ad demonstrates a broader shift in the criteria for the optimal candidate for that specific occupation at that firm at that time, and will thus affect the recruitment patterns of all comparable candidates, regardless of whether they respond to that specific ad or another one published very near in time. Furthermore, the firm may have published an ad for the same occupation in several different places during the same time period and hired more than one candidate, resulting in duplicate matches in the final data.

For each version of the model, I repeat the same exercise but now redefine the digitalisation rate such that the most common technology for each occupation and year is not considered. The idea is to capture the effect of digital skill requirements that might differ somewhat from what is common for that specific occupation during that time period. The motivation is that some digital skill requirements might become increasingly commonplace over time and should therefore not significantly change the demographics among new hires. I expect larger estimates for these models since they capture variation in digital skill requirements that differ more from their non-digital or more commonplace digital counterparts. [7]

4.1 Within-Occupation Variation

To estimate how digital skill requirements affect the share of women among new hires, considering only within-occupation variation, I estimate the following model:

$$\text{Share of women hired}_{ecot} = \gamma \%D_{ecot} + X_{ecot}\beta + \phi_o + \theta_t + \varepsilon_{ecot} \qquad (1)$$

Where $\%D_{ecot}$ is the share of ads that include digital skill requirements of an employer x county interaction, ec, for a specific occupation, o, during year t. X_{ecot} is a vector of controls, such as the employer x county size and the employer x county x occupation size in the previous year. Larger firms may be more prone to introduce new technology and occupations that represent

[7] Table A11 lists the most common technology for each occupation and year.

large shares of workers within firms or workplaces may be more exposed to new technology. The employer x county interaction is used as a proxy for the workplace throughout the analysis due to issues with the workplace indicator in the data. I also includes, for some specifications, the share of women among all employees for that specific workplace and occupation, in the previous period. Fixed effects, such as occupation- and year-fixed effects, are added sequentially. Since the working data do not sample the entire working population and are instead limited to particular occupations, I include occupation fixed effects, ϕ_o, in all iterations of the specification to ensure that only within-occupation variances are considered. Previous work has shown that most changes in tasks and skills, as a result of technological advancements, occur within occupations (Atalay et al., 2020). Year fixed effects, θ_t, are included to capture the general increase in digital skill requirements over all occupations and firms. Standard errors are clustered at the Employer x County level.

All estimations are repeated using the age of new hires as the dependent variable, divided into three age groups. This is to expand the analysis to include important age and gender interactions in the labor market. Young workers include everyone below or equal to 35 years of age, mid-age workers include everyone aged 36 to 50 and older workers include everyone above 50 years of age.

For this group of estimations, one could suspect that higher shares of vacancy ads with digital skill are associated with a lower share of female new hires, considering the stereotypes and educational background distributions discussed earlier. For similar reasons, I also suspect that more digital skill requirements more often lead to the employment of younger workers. In addition, it could also be the case that employers that include digital skill requirements to a larger extent tend to hire relatively fewer women compared to other employers, even for female-dominated occupations.

These models identify the causal relationship between digital skill requirements and employee characteristics, assuming no other gender-related or age-related factors that co-vary with stated digital skill requirements in ads. As discussed above, this is unlikely if firms that deploy digital technologies to a larger extent tend to hire a certain type of worker, perhaps more often male or young. For this reason, the within-occupation results should not be considered causal and the main results stem from estimation models that include employer fixed effects, see Section 4.2. However, the within-occupation models give important descriptive results, where the hiring tendencies of firms that do deploy digital technology are important to capture, rather than control for, as they will determine the trajectory of worker characteristics as digital technologies become more commonplace. These models estimate what happens to occupations overall, including changes that are driven by firm behavior.

4.2 Within-Employer Variation

To estimate how digital skill requirements affect the share of women among new hires, considering only within-employer variation, I estimate the following model:

$$\text{Share of women hired}_{ecot} = \gamma \%D_{ecot} + X_{erot}\beta + \psi_e + \lambda_c + \kappa_{ot} + \varepsilon_{ecot} \quad (2)$$

This is the same model as Model 1, except that each iteration includes employer fixed effects ψ_e. Employer fixed effects are included since employers may differ both in terms of their propensity to hire women and in their technology adoption. Technology adoption has been shown to be correlated with unobservable attributes of the job and variation in skill requirements in ads has been shown to differ more between firms than between different occupations (Athey and Stern, 1998; Deming and Kahn, 2018). I also include county fixed effects, λ_c, to capture characteristics at the local labor market level. Larger cities may be more prone to adopt new technologies and may also attract younger workers to a larger extent. Lastly, I include occupation-specific year fixed effects, κ_{ot}, to control for occupation-specific trends in both digitalisation and demographics among workers. Standard errors are again clustered at the Employer x County level. As with the within-occupation models, I also look at the effect of digital skill requirements on the age of the new recruits, divided into three different age groups.

For this group of estimations, the propensity for a specific employer to include digital skill requirements and hire certain candidates is captured by the employer fixed effects. This allows us to compare, within employer and occupation, how demographics among new recruits respond to the share of vacancy ads that include digital skill requirements. Identifying variation arises from changes in the share of ads that include digital skill requirements from one year to another. Thus, assuming that there are no other gender- or age-related factors that correlate with the inclusion of digital skill requirements in job postings, these models estimate the causal effect of digital technology implementation on the female share and/or age distribution of new recruits.

5 Results

In this section, I present the results of regressing the female share among new hires on the share of ads with digital skill requirements, for both the within-occupation specifications and the within-employer specifications. I also look at the effect of digital skill requirements on age, both pooled and among female workers only.

5.1 Within-Occupation Variation

Table 2 shows that the share of women among new hires does not change significantly when the requirements for digital skills increase. Column 1 includes occupation fixed effects and controls for the employer x county size and the employer x county x occupation size measured in the previous year. Column 2 adds year fixed effects. Column 3 adds the share of women among all employees at the employer x county x occupation level in the previous period that captures the propensity to hire women for that specific occupation at the workplace level. Columns 4-6 mirror columns 1-3 but measure the digitalisation rate whilst excluding the most common technology for each occupation and year. An increase in the digital ad share of one standard deviation (0.38) would increase the female share among new hires by between 0.04 and, at the most, 0.26 percentage points. A very small increase considering the average share of women among new hires is 79% for these occupations.

Table 2. *Within-Occupation Results: Female Share and Digital Skill Requirements*

	Baseline			Most Common Tech. Rem.		
Female Share, New Hires	(1)	(2)	(3)	(4)	(5)	(6)
Digital Ad Share	0.001	0.002	0.003	0.005	0.006	0.007
	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)
Ln(Workplace*Occ. Size)	-0.011***	-0.011***	-0.008***	-0.011***	-0.011***	-0.008***
	(0.002)	(0.002)	(0.001)	(0.002)	(0.002)	(0.001)
Female Share, All Empl. (t-1)			0.291***			0.292***
			(0.012)			(0.012)
R^2	0.10	0.10	0.12	0.10	0.10	0.12
Obs.	32,117	32,117	32,117	32,117	32,117	32,117
Mean Dependent Var.	0.79	0.79	0.79	0.79	0.79	0.79
Workplace Size	✓	✓	✓	✓	✓	✓
Fixed Effects	O	O+Y	O+Y	O	O+Y	O+Y

NOTE.— Each observation is an occupation hired at the Employer x County-level (proxy for workplace), sampled over the years from 2006 to 2016. Columns 4-6 measure the share of ads with digital skill requirements whilst simultaneously disregarding the most common digital skill requirement for each occupation and year. Firm size is divided into buckets 1-9, 10-19, 20-49, 50-249 or 250+ employees and is measured in the previous period. Workplace*Occupation size is measured in the previous period. Standard errors are clustered at the Employer x County level. Singleton observations removed. Missing lagged variables, such as lagged firm size, is set to 0 and controlled for. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 3 presents the relationship between the digital skill requirements and the age distribution of new hires. It shows that the share of young workers (aged 35 or below) increases when the digital skill requirements increase whilst the share of older workers (aged above 50) decreases. The relationship between digital skill requirements and the change in the share of new hires aged 35 or under is sensitive to removing the most common technology when defining the digitalisation rate. The negative relationship between the level of digital skill requirements and the share of new hires over 50 years of age remains robust regardless of digitalisation definition. This result aligns with

previous results and points to a decrease in older workers as occupations become more dependent on digital tools. This result also reassures us that the exercise is actually capturing relevant signs of digital skill requirements in the ad and that the matching of ads to new hires seems to mimic what is expected. A standard deviation increase in the share of digital ads is associated with an increase in younger workers, aged 35 or under, of between 0.3 and 1.3 percentage points, around 0.6 to 3% of the mean. A standard deviation increase in the share of digital ads is also associated with a percentage point decrease of 0.7 to 1.1 of the share of older workers, approximately 4 to 6% of the mean.

Table 4 also presents the results for the age distribution, but includes only female new hires. Again, the relationship between the share of young workers among new hires is sensitive to the digitalisation definition, but the share of older workers is not. A standard deviation increase in the digital ad share is associated with between 0.2 and 1 percentage point increase in the share of younger workers and a 0.7 to 1 percentage point decrease in the share of older workers. Older female workers do not appear to be disproportionately targeted by technology-based age discrimination.

Table 3. *Within-Occupation Results: Age Distribution and Digital Skill Requirements*

	Baseline			Most Common Tech. Rem.		
Share $\leq$ 35, New Hires	(1)	(2)	(3)	(4)	(5)	(6)
Digital Ad Share	0.035***	0.034***	0.021**	0.019*	0.018	0.008
	(0.009)	(0.009)	(0.008)	(0.010)	(0.010)	(0.009)
Ln(Workplace*Occ. Size)	0.021***	0.021***	0.021***	0.021***	0.021***	0.020***
	(0.004)	(0.004)	(0.003)	(0.004)	(0.004)	(0.003)
Share $\leq$ 35, All Empl. (t-1)			0.405***			0.407***
			(0.024)			(0.025)
R^2	0.06	0.06	0.11	0.06	0.06	0.11
Obs.	32,117	32,117	32,117	32,117	32,117	32,117
Mean Dependent Var.	0.47	0.47	0.47	0.47	0.47	0.47
Workplace Size	✓	✓	✓	✓	✓	✓
Fixed Effects	O	O+Y	O+Y	O	O+Y	O+Y
Share 36 – 50, New Hires	(1)	(2)	(3)	(4)	(5)	(6)
Digital Ad Share	-0.006	-0.005	-0.006	0.007	0.008	0.006
	(0.007)	(0.007)	(0.007)	(0.008)	(0.008)	(0.008)
Ln(Workplace*Occ. Size)	-0.012***	-0.012***	-0.012***	-0.012***	-0.011***	-0.012***
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
Share 36 – 50, All Empl. (t-1)			0.191***			0.190***
			(0.018)			(0.018)
R^2	0.03	0.03	0.04	0.03	0.03	0.04
Obs.	32,117	32,117	32,117	32,117	32,117	32,117
Mean Dependent Var.	0.36	0.36	0.36	0.36	0.36	0.36
Workplace Size	✓	✓	✓	✓	✓	✓
Fixed Effects	O	O+Y	O+Y	O	O+Y	O+Y
Share > 50, New Hires	(1)	(2)	(3)	(4)	(5)	(6)
Digital Ad Share	-0.029***	-0.029***	-0.020***	-0.025***	-0.025***	-0.018***
	(0.005)	(0.005)	(0.005)	(0.004)	(0.004)	(0.004)
Ln(Workplace*Occ. Size)	-0.009***	-0.009***	-0.008***	-0.009***	-0.009***	-0.008***
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
Share > 50, All Empl. (t-1)			0.230***			0.230***
			(0.019)			(0.019)
R^2	0.03	0.03	0.05	0.03	0.03	0.05
Obs.	32,117	32,117	32,117	32,117	32,117	32,117
Mean Dependent Var.	0.17	0.17	0.17	0.17	0.17	0.17
Workplace Size	✓	✓	✓	✓	✓	✓
Fixed Effects	O	O+Y	O+Y	O	O+Y	O+Y

NOTE.— Each observation is an occupation hired at the Employer x County-level (proxy for workplace), sampled over the years from 2006 to 2016. Columns 4-6 measure the share of ads with digital skill requirements whilst simultaneously disregarding the most common digital skill requirement for each occupation and year. Firm size is divided into buckets 1-9, 10-19, 20-49, 50-249 or 250+ employees and is measured in the previous period. Workplace x Occupation size is measured in the previous period. Standard errors are clustered at the Employer x County level. Singleton observations removed. Missing lagged variables, such as lagged firm size, is set to 0 and controlled for. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 4. *Within-Occupation Results: Female Age Distribution and Digital Skill Requirements*

	Baseline			Most Common Tech. Rem.		
Fem. Share $\leq$ 35, New Hires	(1)	(2)	(3)	(4)	(5)	(6)
Digital Ad Share	0.026**	0.026**	0.015	0.015	0.015	0.006
	(0.011)	(0.011)	(0.010)	(0.011)	(0.011)	(0.010)
Ln(Workplace*Occ. Size)	0.022***	0.022***	0.024***	0.021***	0.022***	0.024***
	(0.004)	(0.004)	(0.003)	(0.004)	(0.004)	(0.003)
Fem. Share $\leq$ 35,			0.398***			0.399***
All Empl. (t-1)			(0.021)			(0.021)
R^2	0.06	0.06	0.10	0.06	0.06	0.10
Obs.	29,141	29,141	29,141	29,141	29,141	29,141
Mean Dependent Var.	0.46	0.46	0.46	0.46	0.46	0.46
Workplace Size	✓	✓	✓	✓	✓	✓
Fixed Effects	O	O+Y	O+Y	O	O+Y	O+Y
Fem. Share 36 - 50, New Hires	(1)	(2)	(3)	(4)	(5)	(6)
Digital Ad Share	0.000	0.002	0.002	0.009	0.011	0.010
	(0.008)	(0.008)	(0.008)	(0.009)	(0.009)	(0.008)
Ln(Workplace*Occ. Size)	-0.013***	-0.013***	-0.012***	-0.013***	-0.013***	-0.012***
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
Fem. Share 36 – 50,			0.198***			0.198***
All Empl. (t-1)			(0.014)			(0.014)
R^2	0.03	0.03	0.04	0.03	0.03	0.04
Obs.	29,141	29,141	29,141	29,141	29,141	29,141
Mean Dependent Var.	0.37	0.37	0.37	0.37	0.37	0.37
Workplace Size	✓	✓	✓	✓	✓	✓
Fixed Effects	O	O+Y	O+Y	O	O+Y	O+Y
Fem. Share > 50, New Hires	(1)	(2)	(3)	(4)	(5)	(6)
Digital Ad Share	-0.027***	-0.028***	-0.021***	-0.024***	-0.026***	-0.019***
	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.004)
Ln(Workplace*Occ. Size)	-0.009***	-0.009***	-0.008***	-0.009***	-0.009***	-0.008***
	(0.002)	(0.003)	(0.002)	(0.002)	(0.003)	(0.002)
Fem. Share > 50,			0.240***			0.241***
All Empl. (t-1)			(0.018)			(0.018)
R^2	0.03	0.03	0.05	0.03	0.03	0.05
Obs.	29,141	29,141	29,141	29,141	29,141	29,141
Mean Dependent Var.	0.16	0.16	0.16	0.16	0.16	0.16
Workplace Size	✓	✓	✓	✓	✓	✓
Fixed Effects	O	O+Y	O+Y	O	O+Y	O+Y

NOTE.— Each observation is an occupation hired at the Employer x County-level (proxy for workplace), sampled over the years from 2006 to 2016. Columns 4-6 measure the share of ads with digital skill requirements whilst simultaneously disregarding the most common digital skill requirement for each occupation and year. Firm size is divided into buckets 1-9, 10-19, 20-49, 50-249 or 250+ employees and is measured in the previous period. Workplace x Occupation size is measured in the previous period. Standard errors are clustered at the Employer x County level. Singleton observations removed. Missing lagged variables, such as lagged firm size, is set to 0 and controlled for. * p < 0.10, ** p < 0.05, *** p < 0.01

5.2 Within-Employer Variation

It is important to note that none of the above specifications includes employer-fixed effects. Firms adopting software technologies and emphasizing digital skills in their ads are likely to favor younger hires, for example. In this subsection, I present the results from the within-employer variation specifications, which take into account that firms that deploy more technology and require more digital skills from their new hires, might differ in their propensity to hire women and/or workers of different ages. They may also differ in their ability to attract female applicants.

Table 5 includes employer fixed effects in all specifications and considers only within-employer variation in digital skill requirements. The results show that the share of women among new hires increases when the share of vacancy ads with digital skill requirements increases. This is the case in both the baseline model and when the most common technologies are removed when defining the digitalisation rate.[8] This shows that when controlling for employer-specific hiring and technology adoption tendencies, not only do male workers not replace women when digital skill requirements increase, female workers appear more likely to be hired for such positions. The significance level is at 5 to 10% and the results should be interpreted thereafter. In any case, technology-related skill requirements are not associated with a decline in the share of female new hires, but rather that they, if anything, remain the same.

Table 5. *Main Results: Female Share and Digital Skill Requirements*

	Baseline			Most Common Tech. Rem.		
Female Share, New Hires	(1)	(2)	(3)	(4)	(5)	(6)
Digital Ad Share	0.014*	0.014*	0.012	0.018**	0.018**	0.016*
	(0.007)	(0.007)	(0.008)	(0.007)	(0.007)	(0.008)
Ln(Workplace*Occ. Size)	-0.004	-0.004	-0.002	-0.004	-0.003	-0.002
	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)
R^2	0.24	0.24	0.25	0.24	0.24	0.25
Obs.	27,612	27,612	27,612	27,612	27,612	27,612
Mean Dependent Var.	0.78	0.78	0.78	0.78	0.78	0.78
Workplace Size	✓	✓	✓	✓	✓	✓
Fixed Effects	O+Y+E	O+Y+E+C	E+C+O*Y	O+Y+E	O+Y+E+C	E+C+O*Y

NOTE.— Each observation is an occupation hired at the Employer x County-level (proxy for workplace), sampled over the years from 2006 to 2016. Columns 4-6 measure the share of ads with digital skill requirements whilst simultaneously disregarding the most common digital skill requirement for each occupation and year. Firm size is divided into buckets 1-9, 10-19, 20-49, 50-249 or 250+ employees and is measured in the previous period. Workplace*Occupation size is measured in the previous period. Standard errors are clustered at the Employer x County level. Singleton observations removed. Missing lagged variables, such as lagged firm size, is set to 0 and controlled for. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

[8] Tables A13, A14 and A15 show that the results are robust to using a continuous measurement for digital skill requirements.

Column 1 includes occupation-, year-, and employer-fixed effects, column 2 adds county fixed effects, and column 3 adds occupation-specific year fixed effects. Columns 4-6 mirror columns 1-3 but with the alternative digitalisation measurement. A standard deviation increase in the share of digital ads is associated with a 0.5 to 0.7 percentage point increase in the share of women among new hires. As expected, the estimates are larger when the digital ad share does not include the most common technologies.

Table 6 shows, when male and female new hires are pooled, that digital skill requirements are associated with an increase in the share of workers aged 36-50 and a decrease in the share of younger workers. The estimates are, for the most part, not significant, except for columns 4 and 5 in the middle panel, looking at mid-aged workers (10% significance). The share of older workers remains unchanged, with estimates at 0 or very close to 0. This is in contrast to Table 3, which showed an increase in the share of younger workers and a significant decrease in the share of older workers.

Table 7 shows that the relationship between digital skill requirements and the share of older workers is mainly negative, but estimates are small. A standard deviation increase in share digital ads is associated with a 0.5 to 0.8 percentage point increase in the share of female new hires aged 36 to 50 years. This represents around 1.4 to 2.2% of the mean. The significance level is at 10% and the results should be interpreted carefully.

Considering the within-occupation results presented in Tables 2 and 4, my findings point toward a negative between-effect, such that employers that include more digital skill requirements also tend to hire relatively more men and more younger workers overall. In conclusion, the results show that vacancy ads that include digital skill requirements appear to more often leads to the recruitment of a woman aged 36 to 50 compared to men, generally, and women in other age groups.[9][10] In the next subsection, I explore which mechanisms might be driving these results.

[9]Table A16 shows that digital skill requirements have no significant effect on the educational background among new hires.

[10]Tables A17, A18 and A19 show that the overall conclusions still hold when matching on Employer x County x Occupation on a yearly basis, rather than matching on individual ads and new hires.

Table 6. *Age Distribution and Digital Skill Requirements*

	Baseline			Most Common Tech. Rem.		
Share $\leq$ 35, New Hires	(1)	(2)	(3)	(4)	(5)	(6)
Digital Ad Share	-0.007	-0.007	-0.005	-0.017	-0.017	-0.015
	(0.008)	(0.008)	(0.009)	(0.010)	(0.010)	(0.010)
Ln(Workplace*Occ. Size)	0.015***	0.016***	0.015***	0.015***	0.015***	0.015***
	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)
R^2	0.26	0.26	0.27	0.26	0.26	0.27
Obs.	27,612	27,612	27,612	27,612	27,612	27,612
Mean Dependent Var.	0.48	0.48	0.48	0.48	0.48	0.48
Workplace Size	✓	✓	✓	✓	✓	✓
Fixed Effects	O+Y+E	O+Y+E+C	E+C+O*Y	O+Y+E	O+Y+E+C	E+C+O*Y
Share 36 – 50, New Hires	(1)	(2)	(3)	(4)	(5)	(6)
Digital Ad Share	0.007	0.008	0.005	0.016*	0.016*	0.014
	(0.007)	(0.007)	(0.008)	(0.009)	(0.009)	(0.010)
Ln(Workplace*Occ. Size)	-0.010***	-0.010***	-0.010***	-0.010***	-0.010***	-0.010***
	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)
R^2	0.19	0.19	0.20	0.19	0.19	0.20
Obs.	27,612	27,612	27,612	27,612	27,612	27,612
Mean Dependent Var.	0.36	0.36	0.36	0.36	0.36	0.36
Workplace Size	✓	✓	✓	✓	✓	✓
Fixed Effects	O+Y+E	O+Y+E+C	E+C+O*Y	O+Y+E	O+Y+E+C	E+C+O*Y
Share > 50, New Hires	(1)	(2)	(3)	(4)	(5)	(6)
Digital Ad Share	-0.000	-0.001	-0.000	0.001	0.001	0.000
	(0.010)	(0.010)	(0.010)	(0.010)	(0.010)	(0.010)
Ln(Workplace*Occ. Size)	-0.005**	-0.006***	-0.005***	-0.005**	-0.006***	-0.005***
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
R^2	0.20	0.20	0.20	0.20	0.20	0.20
Obs.	27,612	27,612	27,612	27,612	27,612	27,612
Mean Dependent Var.	0.16	0.16	0.16	0.16	0.16	0.16
Workplace Size	✓	✓	✓	✓	✓	✓
Fixed Effects	O+Y+E	O+Y+E+C	E+C+O*Y	O+Y+E	O+Y+E+C	E+C+O*Y

NOTE.— Top Panel: The dependent variable is the share of new hires aged 35 or below. Middle Panel: The dependent variable is the share of new hires aged 36 up to and including 50. Bottom Panel: The dependent variable is the share of workers aged above 50.

Each observation is an occupation hired at the Employer x County-level (proxy for workplace), sampled over the years from 2006 to 2016. Columns 4-6 measure the share of ads with digital skill requirements whilst simultaneously disregarding the most common digital skill requirement for each occupation and year. Firm size is divided into buckets 1-9, 10-19, 20-49, 50-249 or 250+ employees and is measured in the previous period. Workplace*Occupation size is measured in the previous period. Standard errors are clustered at the Employer x County level. Singleton observations removed. Missing lagged variables, such as lagged firm size, is set to 0 and controlled for. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 7. *Female Age Distribution and Digital Skill Requirements*

	Baseline			Most Common Tech. Rem.		
Female Share ≤ 35, **New Hires**	(1)	(2)	(3)	(4)	(5)	(6)
Digital Ad Share	-0.013	-0.013	-0.011	-0.017	-0.017	-0.015
	(0.010)	(0.010)	(0.011)	(0.010)	(0.011)	(0.011)
Ln(Workplace*Occ. Size)	0.015***	0.016***	0.015***	0.015***	0.015***	0.015***
	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)
R^2	0.26	0.26	0.26	0.26	0.26	0.26
Obs.	24,979	24,979	24,979	24,979	24,979	24,979
Mean Dependent Var.	0.47	0.47	0.47	0.47	0.47	0.47
Workplace Size	✓	✓	✓	✓	✓	✓
Fixed Effects	O+Y+E	O+Y+E+C	E+C+O*Y	O+Y+E	O+Y+E+C	E+C+O*Y
Female Share $36-50$, **New Hires**	(1)	(2)	(3)	(4)	(5)	(6)
Digital Ad Share	0.017*	0.017*	0.015*	0.020*	0.021*	0.019*
	(0.008)	(0.008)	(0.009)	(0.010)	(0.010)	(0.011)
Ln(Workplace*Occ. Size)	-0.010***	-0.010***	-0.010***	-0.010***	-0.010***	-0.010***
	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)
R^2	0.19	0.19	0.20	0.19	0.19	0.20
Obs.	24,979	24,979	24,979	24,979	24,979	24,979
Mean Dependent Var.	0.37	0.37	0.37	0.37	0.37	0.37
Workplace Size	✓	✓	✓	✓	✓	✓
Fixed Effects	O+Y+E	O+Y+E+C	E+C+O*Y	O+Y+E	O+Y+E+C	E+C+O*Y
Female Share > 50, **New Hires**	(1)	(2)	(3)	(4)	(5)	(6)
Digital Ad Share	-0.004	-0.004	-0.004	-0.004	-0.004	-0.004
	(0.010)	(0.010)	(0.010)	(0.010)	(0.010)	(0.010)
Ln(Workplace*Occ. Size)	-0.005**	-0.005**	-0.005**	-0.005**	-0.005**	-0.005**
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
R^2	0.20	0.20	0.20	0.20	0.20	0.20
Obs.	24,979	24,979	24,979	24,979	24,979	24,979
Mean Dependent Var.	0.16	0.16	0.16	0.16	0.16	0.16
Workplace Size	✓	✓	✓	✓	✓	✓
Fixed Effects	O+Y+E	O+Y+E+C	E+C+O*Y	O+Y+E	O+Y+E+C	E+C+O*Y

NOTE.— Top Panel: The dependent variable is the share of new hires that are female and aged 35 and below. Middle Panel: The dependent variable is the share of new hires that are female and aged 36 up to and including 50. Bottom Panel: The dependent variable is the share of new hires that are female and aged above 50. Each observation is an occupation hired at the Employer x County-level (proxy for workplace), sampled over the years from 2006 to 2016. Columns 4-6 measure the share of ads with digital skill requirements whilst simultaneously disregarding the most common digital skill requirement for each occupation and year. Firm size is divided into buckets 1-9, 10-19, 20-49, 50-249 or 250+ employees and is measured in the previous period. Workplace*Occupation size is measured in the previous period. Standard errors are clustered at the Employer x County level. Singleton observations removed. Missing lagged variables, such as lagged firm size, is set to 0 and controlled for. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

5.3 Mechanisms

The results show that, when controlling for employer fixed effects, vacancy ads that include digital skill requirements more often lead to the recruitment of a woman, aged 36 to 50 years. In this section, I explore two mechanisms that might be driving these results, previous experience in the occupation and relative wage in previous positions. Both capture the notion that jobs that require more digital or technology-related skills have a higher relative task difficulty and therefore would require more experienced workers (Atalay et al., 2018; Spitz-Oener, 2006).

Previous Experience

Table 8 shows that, among new hires, it is more likely that the new female hires have previously worked in the same occupation in which they are hired. It also shows that the likelihood of previous experience in the said occupation increases with age. This would explain why employers prefer female workers when recruiting for a position with high levels of digital skill requirements. However, it only explains why employers would prefer workers outside of the youngest age group, but not why they prefer "prime age" (36-50) workers above older (50 +) workers, as they do not have the same level of previous experience.

Table 8. *Previous Experience over Gender and Age*

	Share with Previous Experience in Occupation	
	Women	Men
All	.52	.41
Age ≤ 35	.37	.27
Age $36-50$	.62	.52
Age > 50	.78	.66
Obs.	2,638,045	538,546

NOTE.— Share of new hires with previous experience in the occupation that they are hired for, divided by gender and age group.

The top panel of table 9 shows that new hires for ads with digital skill requirements have not, on average, previously worked in the same occupation for which they are being hired. However, the bottom panel of table 9 shows that new female recruits for digital ads have, on average, less previous experience in the occupation for which they are being hired. This makes sense considering that the oldest age group, 50+ years, is not hired to the same extent when the advertisements include digital skill requirements, although they have more experience in that occupation. Therefore, hiring the candidate with the *most* experience in that specific occupation does not appear to be the driving force behind primarily hiring prime-age working women for positions with digital skill requirements. If recruiters only cared about previous experience in the same occupation, the estimates on the share of older workers on the

share of ads with digital skill requirements would have been positive. On the other hand, it could mean that there is a minimum level of experience that is sought after, but it is moderated by the assumption that older workers are less trainable and have shorter employment times.

Table 9. *Previous Experience in Occupation and Digital Skill Requirements*

	Baseline			Most Common Tech. Rem.		
Share w. Prev. Exp., New Hires	(1)	(2)	(3)	(4)	(5)	(6)
Digital Ad Share	-0.013	-0.014	-0.014	-0.013	-0.014	-0.014
	(0.009)	(0.009)	(0.009)	(0.010)	(0.010)	(0.010)
Ln(Workplace*Occ. Size)	-0.011***	-0.012***	-0.013***	-0.011***	-0.012***	-0.013***
	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)
R^2	0.24	0.24	0.25	0.24	0.24	0.25
Obs.	27,612	27,612	27,612	27,612	27,612	27,612
Mean Dependent Var.	0.42	0.42	0.42	0.42	0.42	0.42
Workplace Size	✓	✓	✓	✓	✓	✓
Fixed Effects	O+Y+E	O+Y+E+C	E+C+O*Y	O+Y+E	O+Y+E+C	E+C+O*Y
Female Share w. Prev. Exp., New Hires	(1)	(2)	(3)	(4)	(5)	(6)
Digital Ad Share	-0.018*	-0.019**	-0.019**	-0.018*	-0.018*	-0.018*
	(0.009)	(0.009)	(0.009)	(0.009)	(0.009)	(0.009)
Ln(Workplace*Occ. Size)	-0.006**	-0.007**	-0.009**	-0.006**	-0.007**	-0.009**
	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)
R^2	0.23	0.23	0.24	0.23	0.23	0.24
Obs.	24,979	24,979	24,979	24,979	24,979	24,979
Mean Dependent Var.	0.43	0.43	0.43	0.43	0.43	0.43
Workplace Size	✓	✓	✓	✓	✓	✓
Fixed Effects	O+Y+E	O+Y+E+C	E+C+O*Y	O+Y+E	O+Y+E+C	E+C+O*Y

NOTE.— The dependent variable is the share of new hires that have previous experience in the occupation for which they are being hired. Each observation is an occupation hired at the Employer x County-level (proxy for workplace), sampled over the years from 2006 to 2016. Columns 4-6 measure the share of ads with digital skill requirements whilst simultaneously disregarding the most common digital skill requirement for each occupation and year. Firm size is divided into buckets 1-9, 10-19, 20-49, 50-249 or 250+ employees and is measured in the previous period. Workplace*Occupation size is measured in the previous period. Standard errors are clustered at the Employer x County level. Singleton observations removed. Missing lagged variables, such as lagged firm size, is set to 0 and controlled for.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Previous Wage

In this context, previous experience in the same occupation may not be indicative of the type of skill that is most sought after for a position with relatively higher levels of digital skill requirements. One way to measure relative skill is by relative wage, both within their previous occupations and overall, with the assumption that workers with relatively higher wages are more skilled. For each new hire in my data, I track the previous relative wage (through wage

percentiles) and regress the average wage percentile among new hires on the share of digital ads for a specific occupation hired at a workplace in a given year. Table 10 shows that new hires for positions with digital skill requirements have previously had relatively higher wages, compared to new hires for ads without digital skill requirements. This appears to hold when relative wage is measured within the new hires previous occupation and when measured using the whole sampled labor market. One way to interpret this is that firms hire relatively more skilled candidates when the position specifically requires digital skills. Most age-earnings profiles peak at around age 50 and decline slightly after that, which points to worker productivity peaking before they reach the oldest age group considered here.

Table 10. *Previous Relative Wage and Digital Skill Requirements*

	Baseline			Most Common Tech. Rem.		
Within-Occ. Wage Percentile	(1)	(2)	(3)	(4)	(5)	(6)
Digital Ad Share	0.978**	0.602	0.607	1.496***	1.075***	1.031***
	(0.425)	(0.397)	(0.410)	(0.368)	(0.344)	(0.356)
Ln(Workplace*Occ. Size)	-1.983***	-2.376***	-2.412***	-1.977***	-2.370***	-2.407***
	(0.454)	(0.367)	(0.387)	(0.455)	(0.368)	(0.388)
R^2	0.31	0.33	0.34	0.32	0.33	0.34
Obs.	26,111	26,111	26,111	26,111	26,111	26,111
Mean Dependent Var.	39.82	39.82	39.82	39.82	39.82	39.82
Workplace Size	✓	✓	✓	✓	✓	✓
Fixed Effects	O+Y+E	O+Y+E+C	E+C+O*Y	O+Y+E	O+Y+E+C	E+C+O*Y
Overall Wage Percentile	(1)	(2)	(3)	(4)	(5)	(6)
Digital Ad Share	1.001**	0.691*	0.620	1.351***	1.006**	0.910**
	(0.397)	(0.382)	(0.386)	(0.373)	(0.367)	(0.355)
Ln(Workplace*Occ. Size)	-1.633***	-1.948***	-1.849***	-1.629***	-1.944***	-1.846***
	(0.345)	(0.277)	(0.276)	(0.345)	(0.277)	(0.276)
R^2	0.58	0.59	0.59	0.58	0.59	0.59
Obs.	26,111	26,111	26,111	26,111	26,111	26,111
Mean Dependent Var.	46.26	46.26	46.26	46.26	46.26	46.26
Workplace Size	✓	✓	✓	✓	✓	✓
Fixed Effects	O+Y+E	O+Y+E+C	E+C+O*Y	O+Y+E	O+Y+E+C	E+C+O*Y

NOTE.— Top Panel: The dependent variable is the average wage percentile for a new hire in their previous occupation, where wage percentiles are measured within occupations for the sampled labor market. Bottom Panel: The dependent variable is the average wage percentile for a new hire in their previous occupation, where wage percentiles are measured for the entire sampled labor market. Each observation is an occupation hired at the Employer x County-level (proxy for workplace), sampled over the years from 2006 to 2016. Columns 4-6 measure the share of ads with digital skill requirements whilst simultaneously disregarding the most common digital skill requirement for each occupation and year. Firm size is divided into buckets 1-9, 10-19, 20-49, 50-249 or 250+ employees and is measured in the previous period. Workplace*Occupation size is measured in the previous period. Standard errors are clustered at the Employer x County level. Singleton observations removed. Missing lagged variables, such as lagged firm size, is set to 0 and controlled for. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

6 Conclusion

The 2000s have been characterized by rapid advancements in workplace digitalisation and software innovation at the same time as men continue to be overrepresented among computer-related educational fields and occupations. What happens when occupations, previously low in software exposure, increase in their digital skill requirements? Especially, what happens to such occupations when they are distinctly female dominated? There is reason to believe that male workers will replace female workers as skill requirements for occupations lean more heavily toward mid-to-advanced software technology, with this process going through changes in vacancy ads and subsequent new hires. Similarly, older workers might be overlooked in favor of younger workers who are presumed to be more adept with new technology.

In this paper, I have shown that this is not the case. Although it seems that, when looking at within-occupation variation only, the share of women among new hires remain unchanged when digital skill requirements increase, turning to within-employer variation shows an increase in the share of women among new hires. Specifically, women in prime working age between 36 and 50 years. When digital skill requirements increase, women appear more sought after than men for these occupations. As mentioned previously, I am unable to disentangle demand and supply effects. We do not know whether the applicant pool has changed significantly over time; perhaps the female share of applicants has decreased, but firms continue to hire women to the same or higher extent.

My results are congruent with previous work that shows that increased technology adoption also increases the skill-level of tasks and non-routine analytical tasks (Atalay et al., 2018; Spitz-Oener, 2006). This increases the demand for more experienced workers, which in these female-dominated occupations will primarily be women. Ads with digital skill requirements do not appear to result in hiring candidates with previous experience in the specific occupation. Instead, employers appear to prefer candidates with higher relative previous wages, indicating that they prefer more skilled and/or productive workers when recruiting for positions that are more reliant on digital tools.

The fact that female-dominated occupations remain as such, even though digital skill requirements have increased in most occupations studied here, aligns with previous work. This research shows that the majority of changes in women's tasks/skills come from within-occupation cells, rather than changes in the gender distribution over occupations (Black and Spitz-Oener, 2010). It shows that women and men do not need to switch occupations for large aggregate changes, such as changes in the gender wage gap, to occur. It could also be the case that gender norms regarding who should inhabit certain positions and occupations are stronger than changes in skill demands. Rather than reimagining who can perform certain jobs, workers adapt to the new skills required of them. Men may be reluctant to apply to female-dominated occu-

pations due to social norms, even though they have become more dependent on digital tools. At this time, it does not appear that increased technology adoption makes a noticeable dent in the occupational gender segregation. When increased reliance on digital tools requires workers with previous experience, it will likely exacerbate already existing segregation.

Lastly, although there has been a noticeable increase in digital skill requirements, it is primarily increases in 'lower level' and less open-ended technology, such as accounting or resource planning software. As more advanced digital tools become more available, it will be informative to revisit this analysis and see if, for example, AI tools affect the gender composition in different ways.

References

Acemoglu, Daron, David Autor, Jonathon Hazell, and Pascual Restrepo (2020). "AI and jobs: Evidence from online vacancies". Tech. rep. National Bureau of Economic Research.

Atalay, Enghin, Phai Phongthiengtham, Sebastian Sotelo, and Daniel Tannenbaum (2018). "New technologies and the labor market". *Journal of Monetary Economics* 97, 48–67.

Atalay, Enghin, Phai Phongthiengtham, Sebastian Sotelo, and Daniel Tannenbaum (2020). "The evolution of work in the United States". *American Economic Journal: Applied Economics* 12.2, 1–34.

Athey, Susan and Scott Stern (1998). *An empirical framework for testing theories about complimentarity in organizational design*.

Autor, David H, Frank Levy, and Richard J Murnane (2003). "The skill content of recent technological change: An empirical exploration". *The Quarterly journal of economics* 118.4, 1279–1333.

Autor, David H. (2013). "The "task approach" to labor markets : an overview". *Journal for Labour Market Research* 46.3, 185–199.

Barth, Erling, James C Davis, Richard B Freeman, and Kristina McElheran (2020). "Twisting the Demand Curve: Digitalization and the Older Workforce". Tech. rep. National Bureau of Economic Research.

Beaudry, Paul and Ethan Lewis (2014). "Do male-female wage differentials reflect differences in the return to skill? Cross-city evidence from 1980-2000". *American Economic Journal: Applied Economics* 6.2, 178–94.

Bertrand, Marianne (2011). "New perspectives on gender". *Handbook of labor economics*. Vol. 4. Elsevier, 1543–1590.

Black, Sandra E and Alexandra Spitz-Oener (2010). "Explaining women's success: Technological change and the skill content of women's work". *The Review of Economics and Statistics* 92.1, 187–194.

Cortes, Guido Matias, Nir Jaimovich, and Henry E Siu (2018). "The "end of men" and rise of women in the high-skilled labor market". Tech. rep. National Bureau of Economic Research.

Deming, David and Lisa B Kahn (2018). "Skill requirements across firms and labor markets: Evidence from job postings for professionals". *Journal of Labor Economics* 36.S1, S337–S369.

Goldin, Claudia (2014). "A grand gender convergence: Its last chapter". *American economic review* 104.4, 1091–1119.

Goldin, Claudia and Lawrence F Katz (2016). "A most egalitarian profession: Pharmacy and the evolution of a family-friendly occupation". *Journal of Labor Economics* 34.3, 705–746.

Hershbein, Brad and Lisa B Kahn (2018). “Do recessions accelerate routine-biased technological change? Evidence from vacancy postings”. *American Economic Review* 108.7, 1737–72.

Jensen, Mathias Fjællegaard (2020). “Gender Differences in Returns to Skills”. Tech. rep. Working Paper.

Kettemann, Andreas, Andreas I Mueller, and Josef Zweimüller (2018). “Vacancy durations and entry wages: evidence from linked vacancy-employer-employee data”. *NBER Working Paper* w25118.

Kuhn, Peter and Kailing Shen (2013). “Gender discrimination in job ads: Evidence from China”. *The Quarterly Journal of Economics* 128.1, 287–336.

Lindley, Joanne (2012). “The gender dimension of technical change and the role of task inputs”. *Labour Economics* 19.4, 516–526.

SCB (2017). “Arbetskraftsundersökningarna”. Tech. rep. Statistiska Centralbyrån.

Schleife, Katrin (2006). “Computer use and employment status of older workers—an analysis based on individual data”. *Labour* 20.2, 325–348.

Spitz-Oener, Alexandra (2006). “Technical change, job tasks, and rising educational demands: Looking outside the wage structure”. *Journal of labor economics* 24.2, 235–270.

Van Borm, Hannah, Ian Burn, and Stijn Baert (2021). “What does a job candidate’s age signal to employers?” *Labour Economics* 71, 102003.

Zolas, Nikolas, Zachary Kroff, Erik Brynjolfsson, Kristina McElheran, David N Beede, Cathy Buffington, Nathan Goldschlag, Lucia Foster, and Emin Dinlersoz (2021). “Advanced Technologies Adoption and Use by US Firms: Evidence from the Annual Business Survey”. Tech. rep. National Bureau of Economic Research.

Appendix A: Additional figures and tables

Table A1. *Top 40 Occupations - All Sample Groups*

Vacancy Ads		All Employees		New Hires	
Occupation	**% of Ads**	**Occupation**	**% of Empl.**	**Occupation**	**% of Empl.**
Sales Associate	5.5%	Assistant Nurse	4.7%	Caretaker	5.4%
Caretaker	5.0%	Caretaker	4.5%	Shop Assistant	3.4%
Preschool Teacher	4.8%	Teacher Assistant	2.7%	Assistant Nurse	2.7%
Sales (Call center)	4.7%	Shop Assistant	2.5%	Teacher Assistant	2.6%
Element. School Teacher	3.5%	Element. School Teacher	2.4%	Sales Associate	2.3%
Street Vendor	3.4%	Preschool Teacher	2.3%	Restaurant Assistant	2.3%
Restaurant Assistant	2.3%	Sales Associate	2.2%	Janitor	2.1%
Chef	2.2%	System Developer	2.0%	Shop Assistant, Food	2.1%
Social Secretary	2.1%	Office Assistant	1.8%	Office Assistant	2.1%
System Developer	2.1%	Janitor	1.8%	Element. School Teacher	1.8%
Civil Engineer	2.0%	Shop Assistant, Food	1.7%	System Developer	1.8%
Shop Assistant	2.0%	Restaurant Assistant	1.5%	Warehouse Worker	1.5%
Janitor	1.9%	Assistant Caretaker	1.5%	Financial Assistant	1.5%
Server	1.8%	Warehouse Worker	1.5%	Admin. Assistant	1.4%
Nurse	1.8%	Admin., Public Sec.	1.4%	Driver, Large Vehicles	1.4%
High School Teacher	1.7%	Admin. Assistant	1.4%	Preschool Teacher	1.4%
Teacher Assistant	1.7%	Driver, Large Vehicles	1.3%	Assistant Caretaker	1.2%
Admin. Assistant	1.4%	Financial Assistant	1.3%	Admin., Public Sec.	1.2%
Financial Assistant	1.2%	Workshop Worker	1.1%	Carpenter	1.1%
Assistant Nurse	1.2%	Carpenter	1.1%	Chef	1.0%
Teacher, Special Ed.	1.0%	Support Technician	1.0%	University Teacher	1.0%
Office Assistant	0.9%	Doctor	1.0%	Janitor, Public Spaces	1.0%
Receptionist	0.9%	Janitor, Public Spaces	1.0%	Server	0.9%
Admin., Public Sec.	0.8%	Nurse, Workpl. Health	0.9%	Support Technician	0.9%
Shop Assistant, Food	0.8%	University Teacher	0.9%	Doctor	0.9%
Doctor	0.8%	Chef	0.9%	Manager, Small Busi.	0.9%
Support Technician	0.8%	High School Teacher	0.9%	Receptionist	0.8%
Human Resources	0.8%	Social Secretary	0.8%	Workshop Worker	0.8%
Warehouse Worker	0.8%	Office Secretary	0.8%	High School Teacher	0.8%
Hairdresser/Beautician	0.8%	Technician/Engineer	0.7%	Journalist/Author	0.7%
Nanny/Maid	0.7%	Clerk, Bank	0.7%	Manager, Large Busi.	0.7%
Office Secretary	0.7%	Mechanic, Vehicles	0.7%	Accountant	0.7%
Mechanical Engineer	0.7%	Mechanical Engineer	0.7%	Accountant	0.7%
Aid Officer, Public Sec.	0.7%	Manager, Small Busi.	0.7%	Nurse, Workpl. Health	0.7%
Driver, Small Vehicles	0.7%	Accountant	0.7%	Social Secretary	0.7%
Accountant	0.7%	Journalist/Author	0.7%	Sales (Call center)	0.7%
Janitor, Public Spaces	0.7%	Civil Engineer	0.7%	Packer, By Hand	0.7%
Driver, Large Vehicles	0.6%	Electrician, Installations	0.7%	Financial Advisor	0.7%
Teacher, Vocational	0.6%	Bus Driver	0.6%	Civil Engineer	0.6%
Host/Doorman	0.6%	Packer, By Hand	0.6%	Assistant Mechanic	0.6%

NOTE.— The 40 most common occupations in the vacancy ad data, employee data, and new hires among employee data, aggregated all years 2006 to 2016. Recruitment firms removed prior to tabulation.

Table A2. *Top 40 Industries - Vacancy Ads and Employee Data*

Vacancy Ads		All Employees	
Industry	**% of Ads**	**Industry**	**% of Empl.**
Education	15.7%	Education	11.4%
Social Work	13.6%	Social Work	10.6%
Healthcare	8.6%	Health Care	8.9%
Health and Social Care	8.0%	Health and Social Care	7.8%
Consult.	6.1%	Retail	6.6%
Retail	6.0%	Wholesale	5.0%
Wholesale	5.2%	Public Admin.	4.4%
Restaurant	5.0%	Specialised Construct.	4.1%
Public Admin.	2.5%	Transportation, Land	2.9%
Specialised Construct.	2.2%	Data Programm./Consulting	2.3%
Office Services	1.9%	Restaurant	2.1%
Data Programm./Consulting	1.8%	Property Management/Cleaning	2.1%
Property Management/Cleaning	1.8%	Manufact., Other Machines	1.9%
Other Consumer Services	1.7%	Technical Consult.	1.9%
Technical Consult.	1.7%	Construct.	1.8%
Hotel Business	1.5%	Manufact., Metal Goods	1.8%
Transportation, Land	1.4%	Manufact., Vehicles	1.8%
Advertising, Market Research	1.4%	Trade and Repair, Vehicles	1.6%
Advocacy, Religious Activities	1.3%	Consult.	1.6%
Trade and Repair, Vehicles	1.1%	Financial Services excl. Insurance	1.6%
Sports and Entertainment	1.0%	Advocacy, Religious Activities	1.6%
Construct.	0.9%	Real Estate Operations	1.5%
Real Estate Operations	0.8%	Food Production	1.3%
Telecommunication	0.8%	Warehousing	1.2%
Manufact., Metal Goods	0.7%	Consult., Law and Finance	1.0%
Other ventures	0.7%	Post and Delivery	0.9%
Security and Surveillance	0.7%	Hotel Business	0.9%
Publishing	0.6%	Manufact., Computers/Electronics	0.9%
Consult., Law and Finance	0.6%	Manufact., Paper/Related Goods	0.9%
Film, Video and TV	0.6%	Manufact., Steel/Metal	0.8%
Support Service, Finance and Insurance	0.6%	Manufact., Wood excl. Furniture	0.8%
Manufact., Other Machines	0.6%	Publishing	0.8%
Food Production	0.5%	Sports and Entertainment	0.8%
Warehousing	0.5%	Manufact., Electrical Appliances	0.7%
Agriculture	0.4%	Advertising, Market Research	0.7%
Post and Delivery	0.4%	Office Services	0.6%
Manufact., Vehicles	0.4%	Telecommunication	0.6%
Information Services	0.3%	Supply of Electr., Gas	0.6%
Supply of Electr., Gas	0.3%	Manufact., Rubber/Plastic	0.6%
Financial Services excl. Insurance	0.3%	Insurance	0.6%

NOTE.— The 40 most common industries in the vacancy ad and employee data, aggregated all years 2006 to 2016. Recruitment firms removed prior to tabulation.

Table A3. *Technology List: Accountants - Self-Sourced*

Accounting Technologies		
Access	Filemaker	Norstedt
Adramatch	Fortnox	Norstedts
Agda	Garp	Office
Agresso	Heroma	Office-paketet
AS400	Hogia	Office-programmen
Ataio	Hogias	Officepaketet
Ax	Hoogia	Officepaketets
Axapta	Hyperion	OLF
Naan	IF	Oracle
Basword	IFS	Orfi
Briljant	IFS Applications	Palette
Brilliant	Invoice Manager	Pyramid
Business Objects	Iscala	Qlickview
CDI	Jeeves	Qlikview
Centsoft	Kontek	Raindance
Coda	Lön 300	Real
Cognos	Lawson	Revenue
Comers	M3	Reveny
Concord XAC	Maconomy	SAP
Contempus	Mamut	SAP r/3
Crona	Marathon	SAP/r3
Data Warehouse	Microsoft Dynamics AX	Scala
Dynamics	Microsoft NAV	Shared Service
Dynamics AX	Microsoft Office	Sof one
Edison	Monitor	Softone
Edison Byrå	Movex	SPCS
EKO	MS Excel	Summarum
Excel	MS Office	Swelog
Excelfärdigheter	MS Powerpoint	TIL
Excelkunskaper	NAV	Twin
FDT	Navision	Visma
Fenistra	Nordstedt	Visma/SPCS
FI/CO	Nordstedts	XOR

NOTE.— Self-sourced list of technologies for accountants. Sourced from 500 individual vacancy ads and combined.

Table A4. *Digital Skill Requirements*

Basic computer skills		
AbiWord	FileMaker Pro	MS Publisher
Adobe	FileMaker	MS Word
Acrobat	GNOME	MS Works
Adobe Acrobat	GNOME Office	MS Access
Apple iWork MAC	Google Docs	Excel
Computer Aided Translation	IBM Lotus	Office
CAT-verktyg	KOffice	PowerPoint
CAT	Linux Office	Power Point

Continued on next page

Table A4 – continued from previous page

Corel	Movie Magic	OpenOffice
WordPerfect	MS Access	StarOffice
Corel WordPerfect	MS Excel	Think Free
Corel WordPerfect Office	MS Office	XyWrite
ECDL	MS Power Point	Zoho Office Suite
European Comp. Driving Lic.	MS Project	Zoho
Accounting Software		
24Seven Office	Fortnox	Dynamics NAV
24Seven	Garp	Monitor
Adeko	Gourmet	Movex
Agresso	Hansa	Norstedts Compact
Akelius skatt	Hogia Art	Norstedts
Amanda	Hogia skatt	Norstedts skatt
ASW	Hogia	Oracle
BAAN	IBS	Orion
Balans	IFS	Peoplesoft
Bokföring 97	IFS App.s	Prosit/Raindance
Bravo	Intentia	Pyramid
Briljant	JD Edwards	REBAS
Caesar	Jeeves	Redovisning/2
Coda	MacVonk	R/3
ComActivity	Magenta	SAP
Compiere	Mamut	SAP R/3
Concorde	Mapics	Scala
CROM	Marathon	SoftOne
Devis	MicroMore	Visma
Eko	Microsoft Axapta	Visma SPCS
Entré Pro	Axapta	XOR
FDT	Dynamics	XOR Compact
Formula	Dynamics AX	Å-data
Application Platforms		
Amazon Web Services	IBM Domino	SAP
AWS	IBM WebSphere	Sun Java Enterpr. System
Kafka	IIS	Sun Java
Tomcat	JBoss App. Server	Weblogic
Apache Kafka	Jboss	Weblogic Server
Apache Tomcat	SharePoint	Azure
Apache	Office SharePoint Server	Windows Azure
Docker	MOSS	Windows Server
Glassfish	Oracle JCAPS	
Glassfish Server	JCAPS	
Databases		

Continued on next page

Table A4 – continued from previous page

Access	MySQL	SAP Sybase
DB2	NoSQL	SQL-Base
FileMaker	Oracle	SQL-Server
Informix	PostgreSQL	SQL
Ingres	Progress	Lotus
Lotus Approach	Relationsdatabaser	SAP
Integration Platforms		
Citrix	JEMS	Software AG Intergr. Suite
IBM Intergr. Bus	Biztalk	WebMethods
IBM Websphere MQ	Microsoft Biztalk	Software AG WebMethods
Jboss	Oracle Fusion Middleware	SeeBeyond
JBoss Enterpr. Middleware Suite	Oracle Service Bus	Sun SeeBeyond eGate Intergr.
Operating Systems		
Android	OS/400	VMS
Dos	OSE	VSE
GNU	pSOS	VxWorks
IOS	RTOS	Windows för mobila enh.
IRIX	Solaris	Windows Server
JCL	Symbian OS	Windows
Linux	Ubuntu	
Mac OS X	UNIX	
MVS	VM/CMS	
Programming Tools		
Active Directory	InstallShield	Rational Functional Tester
AdminStudio	IntelliJ	Rational Manual Tester
Dreamweaver	Java Builder	Rational Perform. Tester
Adobe Dreamweaver	JBuilder	Rational RequisitePro
GoLive	Jenkins	Rational Rose
Adobe GoLive	JIRA	Rational SoDA
Agila arbetsmetoder	Kanban	Rational Software Architect
Android studio	Microsoft Autopilot	Rational Software Modeler
Apache Ant	Microsoft Expression Web	SCCM
Apache Maven	Microsoft Intune	Scrum
Apache Struts	Autopilot	Sed
Appium	Expression Web	Selenium
Aptana Studio	Intune	Software Virtual. Solution
Bamboo	Mono	SVS
BBedit	Motif	SQLite
Bean Machine	MS Team Found. Server	Teamcity
Bitbucket	NetBeans	Testlink
Bower	Netobjects Fusion	Testrail

Continued on next page

Table A4 – continued from previous page

C++	Node	Typo3
C++ Builder	Oracle Forms	Unit tester
CICS	Oracle Reports	Visual C++
Cold Fusion	Pagemill	Visual InterDev
Delphi	Pagespinner	Visual Page
Eclipse	PHP on TRAX	Visual Source Safe
Emacs	Pico	Visual Studio.NET
FrontPage	RAD	Webedit
Git	Rational App. Developer	Windows Powershell
Gradle	Rational ClearCase	Wise Installation Express
Grunt	Rational ClearQuest	Xcode
Homesite	ClearCase	
Hudson	ClearQuest	

Programming Languages		
Abap	HDML	R
ActionScript	HTML	React native
Ada	Java Enterpr. Edition	REXX
ADO.NET	JME	RPG
AJAX	Java Micro Edition	Ruby on Rails
APL	JSP	Ruby
Apple-script	Java Server Pages	Scala
ASP	Java Standard Edition	Scheme
ASP.NET	Java	SGML
Assembler	JavaScript	Shell
AWK	JCL	Smalltalk
Bash	JCL-script	SQL
BASIC	Kotlin	Swift
C#	LabView	Synon/2E
C++	Lingo	Tcl/Tk
C	Lisp	TypeScript
Centura	Lotus	UML
Clojure	Lotus-script	Unix-script
Clojurescript	ML	VB-script
Cobol	Iogo	VBA
ColdFusion	MS logo	Verilog
CSS	MXML	VHDL
CUDA	Natural Program. Language	Visual Basic
D	Objective-C	Visual Basic.NET
Datalog	OpenCL	VRML
Delphi	OpenGL	WML
DHTML	Pascal	XAML
Eiffel	Perl	XHTML
Erlang	Perl-script	XML

Continued on next page

Table A4 – continued from previous page

F#.NET	PHP	XSD
Fortran	PL/1	XSL
xBase	Prolog	XSLT
FoxPro	Python	XUL
Haskell	Qbasic	
Frameworks		
Eclipselink	Hibernate	Toplink
JavaBeans	JPA	
Enterpr. JavaBeans	OpenJPA	
Networks		
Ethernet	LAN	Novell Netware
Token Ring 6 VPN	WAN	
Quality Assessment Systems		
cGXP	GCP	ISO 14000
GXP	GLP	OHSAS 18001
GMP	ISO 9000	EMAS
Routing Protocol		
BGP	RIP	EIGRP
EIGRP	BGP	OSPF
OSPF		
Publication Tools		
Alfresco	Oracle	Siteshaper
Drupal	Polopoly	Sitevision
EPiServer	Publech	Typo3
Escenic	RedDot	vBulletin
Joomla	SharePoint	WordPress
Lemoon	Sitecore	WebSphere
Litium Studio	SiteDirect	
Communications Protocol		
Bluetooth-teknik	Local Interconnect Network	RFID-teknik
CAN	MOST	SNA
FlexRay	NetBEUI	SOAP
HTTP	NetBIOS	TCP/IP
HTTPS	NFC	X.25
IP-telefoniteknik	COM	X.400
IPX	Object Component Model	X.500
IR-teknik	PCL/5	XML Web Services
LIN	REST	

Continued on next page

Table A4 – continued from previous page

Digital Creation		
Toon Boom Studio	Maxon Cinema 4D	Autodesk Combustion
Autodesk 3ds Max	McNeel Rhino 3D	eyeon Fusion
Autodesk Alias	Newtek Lightwave 3D	Motion
Autodesk Maya	Smith Micro Anime	The Foundry Nuke
Autodesk MotionBuilder	Smith Micro Poser	Adobe Premiere
Autodesk Mudbox	ZBrush	Final Cut Pro
Autodesk Softimage	Adobe Captivate	Autodesk Smoke
Blender	Adobe Lightroom	Avid Media Composer
caligari trueSpace	Articulate Storyline 360	DaVinci Resolve
Carrara	Autodesk Flint	Lightworks
Fusion 360	Autodesk Flame	Media 100
Houdini	Autodesk Inferno	SpeedEDIT
Luxology Modo	Adobe After Effects	Vegas
MAXON BodyPaint 3D	Autodesk Cleaner	

Table A5. *Digitalisation Rate - All Years and Occupations*

	2006	2007	2008	2009	2010	2011	2012	2013	2014	2015	2016	**Total**
Controller												
% Digital	.3	.4	.47	.47	.51	.54	.55	.58	.55	.65	.63	.54
% Digital (Com. Rem.)	.23	.28	.34	.3	.34	.4	.42	.45	.42	.46	.47	.4
Human Resources												
% Digital	.09	.16	.16	.12	.12	.15	.2	.18	.16	.17	.18	.16
% Digital (Com. Rem.)	.07	.14	.14	.11	.11	.13	.18	.16	.14	.15	.16	.14
PR Consultant												
% Digital	.16	.18	.22	.22	.23	.28	.3	.27	.32	.37	.44	.3
% Digital (Com. Rem.)	.15	.15	.19	.2	.21	.25	.22	.24	.29	.35	.41	.26
Administrator, Public Sector												
% Digital	.12	.13	.1	.13	.14	.15	.16	.19	.2	.18	.19	.16
% Digital (Com. Rem.)	.1	.11	.08	.11	.12	.13	.13	.14	.14	.12	.14	.13
Social Secretary												
% Digital	.02	.03	.03	.04	.04	.05	.05	.05	.05	.05	.06	.05
% Digital (Com. Rem.)	.01	.02	.02	.03	.03	.04	.03	.03	.04	.03	.05	.03
Administrative Assistant												
% Digital	.13	.22	.28	.22	.21	.26	.26	.23	.23	.26	.28	.25
% Digital (Com. Rem.)	.1	.18	.24	.18	.18	.22	.23	.2	.2	.23	.25	.21
Accountant												
% Digital	.4	.45	.45	.43	.5	.5	.5	.48	.48	.5	.53	.48
% Digital (Com. Rem.)	.32	.35	.34	.34	.39	.39	.42	.4	.4	.41	.43	.39
Aid Officer, Public Sector												
% Digital	.01	.03	.02	.03	.04	.04	.03	.03	.03	.03	.04	.03
% Digital (Com. Rem.)	.01	.02	.01	.02	.03	.03	.02	.02	.02	.02	.04	.02
Office Secretary												
% Digital	.15	.18	.15	.19	.15	.12	.13	.1	.11	.11	.12	.13
% Digital (Com. Rem.)	.14	.15	.13	.12	.11	.1	.11	.09	.1	.1	.1	.11
Financial Assistant												
% Digital	.32	.39	.44	.38	.42	.46	.47	.45	.44	.44	.47	.43
% Digital (Com. Rem.)	.26	.3	.33	.31	.33	.36	.4	.39	.37	.38	.4	.36
Office Assistant												
% Digital	.16	.22	.21	.15	.19	.23	.23	.22	.21	.25	.28	.21
% Digital (Com. Rem.)	.12	.18	.17	.11	.15	.19	.2	.19	.18	.2	.23	.18
Receptionist												
% Digital	.12	.12	.11	.08	.11	.11	.12	.09	.11	.14	.16	.12
% Digital (Com. Rem.)	.08	.09	.08	.06	.08	.08	.11	.07	.08	.09	.09	.09
Total												
% Digital	.15	.2	.22	.2	.21	.23	.23	.21	.2	.2	.22	.21
% Digital (Com. Rem.)	.12	.16	.17	.15	.17	.18	.19	.17	.17	.17	.18	.17
Obs.	21,065	23,225	22,085	16,766	25,717	30,806	32,218	32,943	36,598	49,148	58,835	349,406

NOTE.— Digitalisation rate for all occupations, all years. Includes both the baseline definition as well as the more restrictive definition where the most common technologies are not considered.

Table A6. *Hiring Window - Frequency Table*

Hiring window Months	ExO	ExCxO	ExMxO
0	29.04%	29.47%	28.65%
1	42.37%	42.60%	41.80%
2	55.32%	55.54%	55.10%
3	66.92%	67.29%	67.10%
4	76.49%	76.81%	77.04%
5	84.24%	84.87%	85.08%
6	90.06%	90.56%	90.81%
7	94.97%	95.40%	95.42%
8	97.90%	98.25%	98.18%
9	99.46%	99.53%	99.51%
10	99.99%	99.99%	99.99%
11	100.00%	100.00%	100.00%
Obs.	44,996	42,281	41,232

NOTE.— Sample is limited to unique matches based on Employer x Occupation, Employer x County x Occupation and Employer x Municipality x Occupation, respectively. All matches with start date prior ad posting is removed before tabulation.

Table A7. *Matching rates - Vacancy and New Hires*

	ExO		ExCxO		ExMxO	
	Ads	Hires	Ads	Hires	Ads	Hires
3 Month Hiring Window						
Controller	.19	.14	.14	.04	.12	.04
Human Resources	.38	.28	.32	.22	.3	.2
PR Consultant	.11	.04	.08	.03	.08	.03
Administrator, Public Sector	.72	.5	.63	.38	.61	.34
Social Secretary	.71	.73	.69	.68	.67	.64
Administrative Assistant	.31	.23	.26	.16	.24	.14
Accountant	.14	.09	.1	.04	.09	.03
Aid Officer, Public Sector	.53	.5	.49	.42	.46	.36
Office Secretary	.68	.47	.66	.44	.55	.34
Financial Assistant	.22	.09	.18	.07	.17	.06
Office Assistant	.22	.1	.18	.07	.17	.05
Receptionist	.26	.18	.21	.12	.19	.09
Total	.42	.27	.38	.22	.35	.19
6 Month Hiring Window						
Controller	.24	.16	.18	.06	.16	.05
Human Resources	.44	.3	.38	.25	.36	.23
PR Consultant	.13	.04	.1	.03	.09	.03
Administrator, Public Sector	.74	.52	.67	.4	.64	.36
Social Secretary	.73	.74	.71	.69	.69	.65
Administrative Assistant	.34	.24	.29	.17	.26	.15
Accountant	.17	.1	.12	.05	.11	.04
Aid Officer, Public Sector	.56	.52	.52	.44	.48	.39
Office Secretary	.69	.48	.67	.45	.58	.35
Financial Assistant	.25	.1	.21	.08	.2	.07
Office Assistant	.25	.12	.21	.08	.19	.06
Receptionist	.28	.19	.24	.13	.21	.11
Total	.45	.29	.4	.23	.38	.2
Obs.	349,407	1,215,251	349,407	1,215,251	349,407	1,215,251

NOTE.— Matching rates for each granular level of matching for each occupation, Employer x Occupation (ExO), Employer x County x Occupation (ExCxO), and Employer x Municipality x Occupation (ExMxO). Match rates are aggregated over all years 2006-2016.

Table A8. *Avg. Number of Duplicate Matches - All Occupations and Years*

	2006	2007	2008	2009	2010	2011	2012	2013	2014	2015	2016	**Total**
Controller												
Ad Dup.	12.79	7.35	11.88	5.27	12.18	38.82	33.98	22.67	19.59	17.95	19.02	20.5
Empl. Dup.	3.19	7.07	4.68	2.37	1.72	3.48	4.9	2.23	2.16	5.33	4.69	4.25
Human Resources												
Ad Dup.	10.48	11.74	16.44	18.39	30.52	13.62	16.22	14.24	13.66	18.41	18.03	17
Empl. Dup.	6.35	6.51	6.35	4.18	5.47	6.71	8.83	7.81	8.91	17.56	19.52	12.04
PR Consultant												
Ad Dup.	4.65	3.53	6.79	3.41	3.3	3.53	6.9	5.02	3.74	3.8	83.54	22.44
Empl. Dup.	2.97	2.27	2.48	1.8	3.44	3.58	3.72	3.65	3.3	2.9	3.85	3.36
Administrator, Public Sector												
Ad Dup.	112.62	134.98	157.44	124.42	66.54	55.38	120.47	124.39	98.57	93.61	120.91	112.35
Empl. Dup.	28.33	18.1	22.32	13.59	11.9	21.06	23.97	20.05	18.26	18.67	30.69	22.44
Social Secretary												
Ad Dup.	114.66	105.46	78.13	124.13	108.55	113.75	127.71	147.59	192.52	209.59	214.25	174.28
Empl. Dup.	55.47	37.87	47.73	34.22	50.93	65.56	80.47	92.71	139.63	188.45	226.95	146.14
Administrative Assistant												
Ad Dup.	19.56	20.97	16.57	24.06	18.62	25.41	33.48	29.11	67.06	58.88	61.38	52.6
Empl. Dup.	5.51	5.42	7.66	5.98	13.3	15.33	11	9.67	23.55	17.97	29.73	20.92
Accountant												
Ad Dup.	6.64	11.31	6.76	4.13	4.21	5.97	6.3	4.53	3.87	5.99	10.19	6.84
Empl. Dup.	4.09	13.4	4.59	3.8	2.94	11.34	10.52	2.4	1.95	2.96	4.21	6.77
Aid Officer, Public Sector												
Ad Dup.	95.78	91.68	76.37	89.25	56.98	76.81	109.26	104.54	66.58	70.92	79.63	82.78
Empl. Dup.	18.43	14.03	15.6	8.06	11.03	14	14.63	18.6	24.64	27.02	31.24	21.3
Office Secretary												
Ad Dup.	202.05	201.91	234.58	225.25	158.13	159.05	188.24	221.51	196.97	214.59	190.18	200.72
Empl. Dup.	45.24	70.31	76.24	44.54	50.89	83.57	91.65	118.63	113.28	156.77	183.76	115.42
Financial Assistant												
Ad Dup.	5.38	5.3	7.13	5.33	5.38	8.96	11.11	10.08	11.07	9.66	11.49	8.92
Empl. Dup.	3.84	34.39	8.64	3.75	3.85	19.54	4.57	7.36	5.13	4.4	8.24	9.59
Office Assistant												
Ad Dup.	81.34	88.66	70.39	39.81	57.59	31.05	57.86	58.32	7	10.55	10.67	68.34
Empl. Dup.	15.54	6.87	2.68	3.9	10.19	3.88	3.49	2.79	2.51	8.55	7.27	9.88
Receptionist												
Ad Dup.	19.31	24.28	32.68	22.88	21.54	23.22	27.61	28.17	21.5	38.95	73.06	36.8
Empl. Dup.	7.11	18.09	19.31	3.93	8.11	10.97	16.13	9.34	6.14	19.68	22.3	15.51
Total												
Ad Dup.	122	134.59	143.16	149.21	102.56	107.48	132.7	155.92	162.76	177.07	176.28	154.09
Empl. Dup.	37.71	39.9	48.04	30.38	37.72	55.57	62.12	78.75	101.58	140.68	167.08	100.96
Obs.	178,367	175,887	168,716	103,750	134,257	202,689	247,093	287,757	372,927	552,755	752,660	3,176,858

NOTE.— Average number of duplicate matches for each occupation and year. New hires are matched with vacancy ads based on employer, county and occupational code using a 6 month hiring window.

Table A9. *Digitalisation Rate - Pre- and Post-Match*

	Pre-Match	Post-Match	Post-Match Unique Obs.
Controller			
Share Digital	.54	.35	.49
Share Digital (Common Rem.)	.40	.26	.37
Human Resources			
Share Digital	.16	.14	.13
Share Digital (Common Rem.)	.14	.12	.12
PR Consultant			
Share Digital	.30	.53	.42
Share Digital (Common Rem.)	.26	.50	.38
Administrator, Public Sector			
Share Digital	.16	.14	.16
Share Digital (Common Rem.)	.13	.12	.13
Social Secretary			
Share Digital	.05	.04	.05
Share Digital (Common Rem.)	.03	.03	.04
Administrative Assistant			
Share Digital	.25	.28	.28
Share Digital (Common Rem.)	.21	.22	.24
Accountant			
Share Digital	.48	.55	.50
Share Digital (Common Rem.)	.39	.49	.43
Aid Officer, Public Sector			
Share Digital	.03	.02	.03
Share Digital (Common Rem.)	.02	.02	.02
Office Secretary			
Share Digital	.13	.13	.11
Share Digital (Common Rem.)	.11	.11	.09
Financial Assistant			
Share Digital	.43	.48	.46
Share Digital (Common Rem.)	.36	.39	.37
Office Assistant			
Share Digital	.21	.18	.23
Share Digital (Common Rem.)	.18	.16	.19
Receptionist			
Share Digital	.12	.14	.13
Share Digital (Common Rem.)	.09	.05	.08
Total			
Share Digital	.21	.10	.15
Share Digital (Common Rem.)	.17	.08	.12
Obs.	349,406	3,176,858	140,958

NOTE.— Digitalisation rate for each occupation, pre- and post-matching. Vacancy ads are matched to new hires based on employer, county and occupation using a 6 month hiring window.

Table A10. *Employee Descriptives - Pre- and Post-Match*

	Pre-Match	Post-Match	Post-Match Unique Obs.
Controller			
Female Share	.57	.59	.59
Avg. Age	40.17	35.68	35.85
Human Resources			
Female Share	.74	.8	.79
Avg. Age	41.39	39.19	40.28
PR Consultant			
Female Share	.52	.52	.55
Avg. Age	38.2	35.08	35.52
Administrator, Public Sector			
Female Share	.6	.64	.64
Avg. Age	43.1	39.63	41.19
Social Secretary			
Female Share	.84	.86	.85
Avg. Age	39.91	35.58	38.39
Administrative Assistant			
Female Share	.74	.76	.77
Avg. Age	40.11	36.1	37.91
Accountant			
Female Share	.76	.78	.79
Avg. Age	42.41	38.59	38.85
Aid Officer, Public Sector			
Female Share	.6	.59	.59
Avg. Age	39.16	36.4	37.76
Office Secretary			
Female Share	.94	.96	.97
Avg. Age	43.5	40.68	43.44
Financial Assistant			
Female Share	.87	.86	.87
Avg. Age	42.05	38.03	39.05
Office Assistant			
Female Share	.73	.76	.77
Avg. Age	38.83	36.73	37.44
Receptionist			
Female Share	.81	.75	.78
Avg. Age	32.8	29.15	30.46
Total			
Female Share	.73	.83	.76
Avg. Age	40.17	37.52	39.09
Obs.	1,215,551	3,176,592	276,977

NOTE.— Female share and average age for each occupation, pre- and post matching with vacancy ads. New hires are matched with vacancy ads based on employer, county and occupational code using a 6 month hiring window.

Table A11. *Most Common Technology - Occupation and Year*

	Controller	Human Resources	PR Consultant	Administrator, Public Sector	Social Secretary	Administrative Assistant
2006	excel	excel	office	excel	rad	excel
2007	excel	excel	excel	excel	rad	excel
2008	excel	excel	office	excel	c	office
2009	excel	office	office	balans	rad	excel
2010	excel	office	office	office	rad	office
2011	excel	office	office	excel	rad	excel
2012	excel	office	excel	balans	c rad	office
2013	excel	excel	office	balans	c	office
2014	excel	office	office	balans	c	office
2015	excel	office	can	balans	rad	office
2016	excel	office	can	balans	visma	office

	Accountant	Aid Officer, Public Sector	Office Secretary	Financial Assistant	Office Assistant	Receptionist
2006	excel	excel	office	excel	excel	office
2007	excel	d	excel	excel	excel	office
2008	excel	c	excel	excel	excel	office
2009	excel	balans	rad	excel	excel	office
2010	excel	rad	excel office	excel	excel	office
2011	excel	office	excel	excel	office	office
2012	excel	balans	office	excel	office	can
2013	excel	c rad	office	excel	office	office
2014	excel	rad	office	excel	excel	office
2015	excel	rad	office	excel	excel	office
2016	excel	visma	excel	excel	excel	office

NOTE.— Most common technology for each occupation and year.

Figure A1. Digitalisation Rate of Occupations - Continuous Measure of Digital Skill Requirements

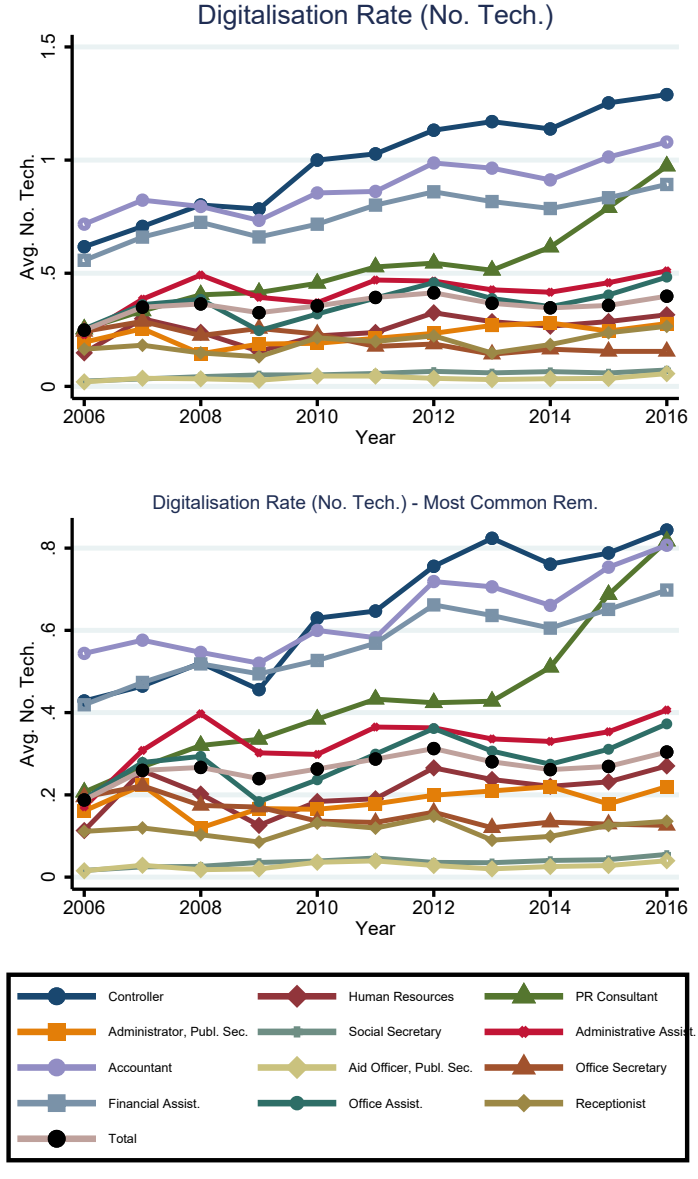

NOTE.—Digitalisation rate over occupations from 2006-2016, prior to matching with new hires. Digitalisation rate measured using a continuous measure of digital skill requirements.

Table A12. *Digitalisation Rate, Avg. No. Technologies - All Occupations and Years*

	2006	2007	2008	2009	2010	2011	2012	2013	2014	2015	2016	**Total**
Controller												
Avg. No.	.62	.71	.8	.78	1	1.03	1.13	1.17	1.14	1.25	1.29	1.06
Avg. No. (Com. Rem.)	.43	.46	.52	.46	.63	.65	.76	.82	.76	.79	.84	.69
Human Resources												
Avg. No.	.15	.3	.24	.16	.22	.24	.33	.29	.27	.29	.32	.27
Avg. No. (Com. Rem.)	.11	.26	.2	.13	.18	.19	.27	.24	.22	.23	.27	.22
PR Consultant												
Avg. No.	.25	.33	.4	.41	.46	.53	.54	.51	.62	.79	.97	.59
Avg. No. (Com. Rem.)	.21	.27	.32	.34	.38	.43	.42	.43	.51	.69	.82	.49
Administrator, Public Sector												
Avg. No.	.2	.25	.14	.19	.19	.21	.23	.27	.28	.24	.28	.24
Avg. No. (Com. Rem.)	.16	.22	.12	.17	.16	.18	.2	.21	.22	.18	.22	.19
Social Secretary												
Avg. No.	.02	.03	.04	.05	.05	.06	.07	.06	.07	.06	.07	.06
Avg. No. (Com. Rem.)	.02	.02	.03	.04	.04	.05	.04	.04	.04	.04	.06	.04
Administrative Assistant												
Avg. No.	.23	.39	.49	.39	.37	.47	.47	.43	.42	.46	.51	.44
Avg. No. (Com. Rem.)	.17	.31	.4	.3	.3	.36	.36	.34	.33	.35	.41	.35
Accountant												
Avg. No.	.72	.82	.8	.73	.85	.86	.99	.96	.91	1.01	1.08	.91
Avg. No. (Com. Rem.)	.54	.58	.55	.52	.6	.58	.72	.71	.66	.75	.81	.65
Aid Officer, Public Sector												
Avg. No.	.02	.04	.03	.03	.05	.05	.04	.03	.03	.04	.06	.04
Avg. No. (Com. Rem.)	.02	.03	.02	.02	.04	.04	.03	.02	.03	.03	.04	.03
Office Secretary												
Avg. No.	.24	.29	.22	.26	.23	.18	.19	.14	.16	.15	.15	.19
Avg. No. (Com. Rem.)	.19	.22	.17	.17	.14	.13	.16	.12	.13	.13	.13	.15
Financial Assistant												
Avg. No.	.56	.66	.73	.66	.72	.8	.86	.82	.79	.83	.89	.78
Avg. No. (Com. Rem.)	.42	.47	.52	.49	.53	.57	.66	.64	.61	.65	.7	.59
Office Assistant												
Avg. No.	.25	.36	.38	.25	.32	.39	.46	.39	.35	.4	.48	.37
Avg. No. (Com. Rem.)	.19	.28	.29	.18	.24	.3	.36	.31	.27	.31	.37	.28
Receptionist												
Avg. No.	.16	.18	.15	.13	.21	.2	.22	.15	.18	.24	.27	.2
Avg. No. (Com. Rem.)	.11	.12	.1	.09	.13	.12	.15	.09	.1	.13	.14	.12
Total												
Avg. No.	.25	.35	.36	.33	.36	.39	.41	.37	.35	.36	.4	.36
Avg. No. (Com. Rem.)	.19	.26	.27	.24	.26	.29	.31	.28	.26	.27	.3	.27
Obs.	21,065	23,225	22,085	16,766	25,717	30,806	32,218	32,943	36,598	49,148	58,835	349,406

NOTE.— Average number of digital skill mentions per occupation and year. Also includes the more restrictive definition where the most common technology for each occupation and year is not considered.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A13. *Robustness - Avg. Number of Digital Skill Requirements: Female Share*

	Baseline			Most Common Tech. Rem.		
Female Share, New Hires	(1)	(2)	(3)	(4)	(5)	(6)
Avg. No. Digital Tech	0.007**	0.007**	0.007*	0.010***	0.010***	0.009**
	(0.003)	(0.003)	(0.004)	(0.003)	(0.003)	(0.004)
Ln(Workplace*Occ. Size)	-0.004	-0.004	-0.002	-0.004	-0.004	-0.002
	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)
R^2	0.24	0.24	0.25	0.24	0.24	0.25
Obs.	27,612	27,612	27,612	27,612	27,612	27,612
Mean Dependent Var.	0.78	0.78	0.78	0.78	0.78	0.78
Workplace Size	✓	✓	✓	✓	✓	✓
Fixed Effects	O+Y+E	O+Y+E+C	E+C+O*Y	O+Y+E	O+Y+E+C	E+C+O*Y

NOTE.—Each observation is an occupation hired at the Employer x County-level (proxy for workplace), sampled over the years from 2006 to 2016. Columns 4-6 measure the share of ads with digital skill requirements whilst simultaneously disregarding the most common digital skill requirement for each occupation and year. Firm size is divided into buckets 1-9, 10-19, 20-49, 50-249 or 250+ employees and is measured in the previous period. Workplace x Occupation size is measured in the previous period. Standard errors are clustered at the Employer x County level. Singleton observations removed. Missing lagged variables, such as lagged firm size, is set to 0 and controlled for.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A14. *Robustness - Avg. Number of Digital Skill Requirements: Age Distribution*

	Baseline			Most Common Tech. Rem.		
Share $\leq$ 35, New Hires	(1)	(2)	(3)	(4)	(5)	(6)
Avg. No. Digital Tech	-0.007*	-0.007*	-0.006	-0.011**	-0.012**	-0.011**
	(0.004)	(0.004)	(0.004)	(0.004)	(0.005)	(0.005)
Ln(Workplace*Occ. Size)	0.015***	0.015***	0.015***	0.015***	0.015***	0.015***
	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)
R^2	0.26	0.26	0.27	0.26	0.26	0.27
Obs.	27,612	27,612	27,612	27,612	27,612	27,612
Mean Dependent Var.	0.48	0.48	0.48	0.48	0.48	0.48
Workplace Size	✓	✓	✓	✓	✓	✓
Fixed Effects	O+Y+E	O+Y+E+C	E+C+O*Y	O+Y+E	O+Y+E+C	E+C+O*Y
Share 36 – 50, New Hires	(1)	(2)	(3)	(4)	(5)	(6)
Avg. No. Digital Tech	0.008**	0.008**	0.008*	0.012**	0.013**	0.012**
	(0.004)	(0.004)	(0.004)	(0.005)	(0.005)	(0.005)
Ln(Workplace*Occ. Size)	-0.010***	-0.010***	-0.010***	-0.010***	-0.010***	-0.010***
	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)
R^2	0.19	0.19	0.20	0.19	0.19	0.20
Obs.	27,612	27,612	27,612	27,612	27,612	27,612
Mean Dependent Var.	0.36	0.36	0.36	0.36	0.36	0.36
Workplace Size	✓	✓	✓	✓	✓	✓
Fixed Effects	O+Y+E	O+Y+E+C	E+C+O*Y	O+Y+E	O+Y+E+C	E+C+O*Y
Share > 50, New Hires	(1)	(2)	(3)	(4)	(5)	(6)
Avg. No. Digital Tech	-0.001	-0.001	-0.001	-0.001	-0.001	-0.001
	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)
Ln(Workplace*Occ. Size)	-0.005***	-0.006***	-0.005***	-0.005***	-0.006***	-0.005***
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
R^2	0.20	0.20	0.20	0.20	0.20	0.20
Obs.	27,612	27,612	27,612	27,612	27,612	27,612
Mean Dependent Var.	0.16	0.16	0.16	0.16	0.16	0.16
Workplace Size	✓	✓	✓	✓	✓	✓
Fixed Effects	O+Y+E	O+Y+E+C	E+C+O*Y	O+Y+E	O+Y+E+C	E+C+O*Y

NOTE.— Top Panel: The dependent variable is the share of new hires aged 35 or below. Middle Panel: The dependent variable is the share of new hires aged 36 up to and including 50. Bottom Panel: The dependent variable is the share of workers aged above 50.

Each observation is an occupation hired at the Employer x County-level (proxy for workplace), sampled over the years from 2006 to 2016. Columns 4-6 measure the share of ads with digital skill requirements whilst simultaneously disregarding the most common digital skill requirement for each occupation and year. Firm size is divided into buckets 1-9, 10-19, 20-49, 50-249 or 250+ employees and is measured in the previous period. Workplace x Occupation size is measured in the previous period. Standard errors are clustered at the Employer x County level. Singleton observations removed. Missing lagged variables, such as lagged firm size, is set to 0 and controlled for.

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

Table A15. *Robustness - Avg. Number of Digital Skill Requirements: Female Age Distribution*

	Baseline			Most Common Tech. Rem.		
Female Share ≤ 35, **New Hires**	(1)	(2)	(3)	(4)	(5)	(6)
Avg. No. Digital Tech	-0.012**	-0.012**	-0.011**	-0.014***	-0.014**	-0.013**
	(0.004)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)
Ln(Workplace*Occ. Size)	0.015***	0.015***	0.015***	0.015***	0.015***	0.015***
	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)
R^2	0.26	0.26	0.26	0.26	0.26	0.26
Obs.	24,979	24,979	24,979	24,979	24,979	24,979
Mean Dependent Var.	0.47	0.47	0.47	0.47	0.47	0.47
Workplace Size	✓	✓	✓	✓	✓	✓
Fixed Effects	O+Y+E	O+Y+E+C	E+C+O*Y	O+Y+E	O+Y+E+C	E+C+O*Y
Female Share $36-50$, **New Hires**	(1)	(2)	(3)	(4)	(5)	(6)
Avg. No. Digital Tech	0.013***	0.014***	0.013***	0.016***	0.016***	0.015***
	(0.004)	(0.004)	(0.004)	(0.005)	(0.005)	(0.005)
Ln(Workplace*Occ. Size)	-0.010***	-0.010***	-0.010***	-0.010***	-0.010***	-0.010***
	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)
R^2	0.19	0.19	0.20	0.19	0.19	0.20
Obs.	24,979	24,979	24,979	24,979	24,979	24,979
Mean Dependent Var.	0.37	0.37	0.37	0.37	0.37	0.37
Workplace Size	✓	✓	✓	✓	✓	✓
Fixed Effects	O+Y+E	O+Y+E+C	E+C+O*Y	O+Y+E	O+Y+E+C	E+C+O*Y
Female Share > 50, **New Hires**	(1)	(2)	(3)	(4)	(5)	(6)
Avg. No. Digital Tech	-0.001	-0.001	-0.001	-0.002	-0.002	-0.002
	(0.004)	(0.004)	(0.004)	(0.005)	(0.005)	(0.005)
Ln(Workplace*Occ. Size)	-0.005**	-0.005**	-0.005**	-0.005**	-0.005**	-0.005**
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
R^2	0.20	0.20	0.20	0.20	0.20	0.20
Obs.	24,979	24,979	24,979	24,979	24,979	24,979
Mean Dependent Var.	0.16	0.16	0.16	0.16	0.16	0.16
Workplace Size	✓	✓	✓	✓	✓	✓
Fixed Effects	O+Y+E	O+Y+E+C	E+C+O*Y	O+Y+E	O+Y+E+C	E+C+O*Y

NOTE.— Top Panel: The dependent variable is the share of new hires that are female and aged 35 and below. Middle Panel: The dependent variable is the share of new hires that are female and aged 36 up to and including 50. Bottom Panel: The dependent variable is the share of new hires that are female and aged above 50. Each observation is an occupation hired at the Employer x County-level (proxy for workplace), sampled over the years from 2006 to 2016. Columns 4-6 measure the share of ads with digital skill requirements whilst simultaneously disregarding the most common digital skill requirement for each occupation and year. Firm size is divided into buckets 1-9, 10-19, 20-49, 50-249 or 250+ employees and is measured in the previous period. Workplace x Occupation size is measured in the previous period. Standard errors are clustered at the Employer x County level. Singleton observations removed. Missing lagged variables, such as lagged firm size, is set to 0 and controlled for.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A16. *Additional Results - Digital Skill Requirements and Educational Background*

	Baseline			Most Common Tech. Rem.		
Share Data Edu., New Hires	(1)	(2)	(3)	(4)	(5)	(6)
Digital Ad Share	0.002	0.002	0.002	0.003	0.003	0.003
	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)
Ln(Workplace*Occ. Size)	0.000	0.000	0.000	0.000	0.000	0.000
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
R^2	0.16	0.16	0.16	0.16	0.16	0.16
Obs.	27,612	27,612	27,612	27,612	27,612	27,612
Mean Dependent Var.	0.01	0.01	0.01	0.01	0.01	0.01
Workplace Size	✓	✓	✓	✓	✓	✓
Fixed Effects	O+Y+E	O+Y+E+C	E+C+O*Y	O+Y+E	O+Y+E+C	E+C+O*Y
Female Share Data Edu., New Hires	(1)	(2)	(3)	(4)	(5)	(6)
Digital Ad Share	0.003	0.003	0.003	0.005	0.005	0.005
	(0.002)	(0.002)	(0.003)	(0.003)	(0.003)	(0.004)
Ln(Workplace*Occ. Size)	0.000	0.000	0.000	0.000	0.000	0.000
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
R^2	0.15	0.15	0.15	0.15	0.15	0.15
Obs.	24,979	24,979	24,979	24,979	24,979	24,979
Mean Dependent Var.	0.01	0.01	0.01	0.01	0.01	0.01
Workplace Size	✓	✓	✓	✓	✓	✓
Fixed Effects	O+Y+E	O+Y+E+C	E+C+O*Y	O+Y+E	O+Y+E+C	E+C+O*Y

NOTE.— Top Panel: The dependent variable is the share of new hires that have an educational background in computer science or related fields. Bottom Panel: The dependent variable is the share of new hires that are both female and have an educational background in computer science or related fields. Each observation is an occupation hired at the Employer x County-level (proxy for workplace), sampled over the years from 2006 to 2016. Columns 4-6 measure the share of ads with digital skill requirements whilst simultaneously disregarding the most common digital skill requirement for each occupation and year. Firm size is divided into buckets 1-9, 10-19, 20-49, 50-249 or 250+ employees and is measured in the previous period. Workplace x Occupation size is measured in the previous period. Standard errors are clustered at the Employer x County level. Singleton observations removed. Missing lagged variables, such as lagged firm size, is set to 0 and controlled for.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A17. *Robustness - Non-Individual Matching: Female Share*

	Baseline			Most Common Tech. Rem.		
Female Share, New Hires	(1)	(2)	(3)	(4)	(5)	(6)
Digital Ad Share	0.010*	0.010*	0.008	0.012**	0.012*	0.011*
	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)
Ln(Workplace*Occ. Size)	-0.007***	-0.007***	-0.006**	-0.007***	-0.007***	-0.006***
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
R^2	0.32	0.32	0.33	0.32	0.32	0.33
Obs.	43,287	43,287	43,287	43,287	43,287	43,287
Mean Dependent Var.	0.77	0.77	0.77	0.77	0.77	0.77
Workplace Size	✓	✓	✓	✓	✓	✓
Fixed Effects	O+Y+E	O+Y+E+C	E+C+O*Y	O+Y+E	O+Y+E+C	E+C+O*Y

NOTE.— Each observation is an occupation hired at the Employer x County-level (proxy for workplace), sampled over the years from 2006 to 2016. Columns 4-6 measure the share of ads with digital skill requirements whilst simultaneously disregarding the most common digital skill requirement for each occupation and year. Firm size is divided into buckets 1-9, 10-19, 20-49, 50-249 or 250+ employees and is measured in the previous period. Workplace x Occupation size is measured in the previous period. Standard errors are clustered at the Employer x County level. Singleton observations removed. Missing lagged variables, such as lagged firm size, is set to 0 and controlled for.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A18. *Robustness - Non-Individual Matching: Age Distribution*

	Baseline			Most Common Tech. Rem.		
Share $\leq$ 35, New Hires	(1)	(2)	(3)	(4)	(5)	(6)
Digital Ad Share	0.000	0.000	0.003	-0.002	-0.003	0.000
	(0.007)	(0.007)	(0.006)	(0.006)	(0.006)	(0.006)
Ln(Workplace*Occ. Size)	0.010***	0.010***	0.010***	0.010***	0.010***	0.010***
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
R^2	0.35	0.35	0.36	0.35	0.35	0.36
Obs.	43,287	43,287	43,287	43,287	43,287	43,287
Mean Dependent Var.	0.41	0.41	0.41	0.41	0.41	0.41
Workplace Size	✓	✓	✓	✓	✓	✓
Fixed Effects	O+Y+E	O+Y+E+C	E+C+O*Y	O+Y+E	O+Y+E+C	E+C+O*Y
Share 36 – 50, New Hires	(1)	(2)	(3)	(4)	(5)	(6)
Digital Ad Share	0.009	0.009	0.006	0.010	0.010	0.008
	(0.005)	(0.005)	(0.006)	(0.006)	(0.006)	(0.006)
Ln(Workplace*Occ. Size)	-0.004**	-0.004**	-0.004**	-0.004**	-0.004**	-0.004**
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
R^2	0.23	0.23	0.24	0.23	0.23	0.24
Obs.	43,287	43,287	43,287	43,287	43,287	43,287
Mean Dependent Var.	0.37	0.37	0.37	0.37	0.37	0.37
Workplace Size	✓	✓	✓	✓	✓	✓
Fixed Effects	O+Y+E	O+Y+E+C	E+C+O*Y	O+Y+E	O+Y+E+C	E+C+O*Y
Share > 50, New Hires	(1)	(2)	(3)	(4)	(5)	(6)
Digital Ad Share	-0.009*	-0.009*	-0.010**	-0.007**	-0.007**	-0.008**
	(0.005)	(0.004)	(0.004)	(0.003)	(0.003)	(0.003)
Ln(Workplace*Occ. Size)	-0.007***	-0.007***	-0.007***	-0.007***	-0.007***	-0.007***
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
R^2	0.28	0.28	0.29	0.28	0.28	0.29
Obs.	43,285	43,285	43,285	43,285	43,285	43,285
Mean Dependent Var.	0.22	0.22	0.22	0.22	0.22	0.22
Workplace Size	✓	✓	✓	✓	✓	✓
Fixed Effects	O+Y+E	O+Y+E+C	E+C+O*Y	O+Y+E	O+Y+E+C	E+C+O*Y

NOTE.— Top Panel: The dependent variable is the share of new hires aged 35 or below. Middle Panel: The dependent variable is the share of new hires aged 36 up to and including 50. Bottom Panel: The dependent variable is the share of workers aged above 50.

Each observation is an occupation hired at the Employer x County-level (proxy for workplace), sampled over the years from 2006 to 2016. Columns 4-6 measure the share of ads with digital skill requirements whilst simultaneously disregarding the most common digital skill requirement for each occupation and year. Firm size is divided into buckets 1-9, 10-19, 20-49, 50-249 or 250+ employees and is measured in the previous period. Workplace x Occupation size is measured in the previous period. Standard errors are clustered at the Employer x County level. Singleton observations removed. Missing lagged variables, such as lagged firm size, is set to 0 and controlled for.

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

Table A19. *Robustness - Non-Individual Matching: Female Age Distribution*

	Baseline			Most Common Tech. Rem.		
Female Share ≤ 35, **New Hires**	(1)	(2)	(3)	(4)	(5)	(6)
Digital Ad Share	0.002	0.001	0.004	0.001	0.001	0.003
	(0.008)	(0.008)	(0.008)	(0.006)	(0.006)	(0.006)
Ln(Workplace*Occ. Size)	0.010***	0.010***	0.010***	0.010***	0.010***	0.010***
	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)
R^2	0.34	0.34	0.34	0.34	0.34	0.34
Obs.	39,917	39,917	39,917	39,917	39,917	39,917
Mean Dependent Var.	0.41	0.41	0.41	0.41	0.41	0.41
Workplace Size	✓	✓	✓	✓	✓	✓
Fixed Effects	O+Y+E	O+Y+E+C	E+C+O*Y	O+Y+E	O+Y+E+C	E+C+O*Y
Female Share $36-50$, **New Hires**	(1)	(2)	(3)	(4)	(5)	(6)
Digital Ad Share	0.007	0.007	0.005	0.007	0.007	0.005
	(0.007)	(0.007)	(0.007)	(0.006)	(0.006)	(0.006)
Ln(Workplace*Occ. Size)	-0.004*	-0.003*	-0.003	-0.004*	-0.003*	-0.003
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
R^2	0.23	0.23	0.23	0.23	0.23	0.23
Obs.	39,917	39,917	39,917	39,917	39,917	39,917
Mean Dependent Var.	0.38	0.38	0.38	0.38	0.38	0.38
Workplace Size	✓	✓	✓	✓	✓	✓
Fixed Effects	O+Y+E	O+Y+E+C	E+C+O*Y	O+Y+E	O+Y+E+C	E+C+O*Y
Female Share > 50, **New Hires**	(1)	(2)	(3)	(4)	(5)	(6)
Digital Ad Share	-0.008**	-0.008**	-0.009**	-0.008**	-0.007**	-0.008**
	(0.004)	(0.004)	(0.003)	(0.004)	(0.004)	(0.003)
Ln(Workplace*Occ. Size)	-0.007**	-0.007**	-0.007**	-0.007**	-0.007**	-0.007**
	(0.002)	(0.002)	(0.003)	(0.002)	(0.002)	(0.003)
R^2	0.26	0.27	0.27	0.26	0.27	0.27
Obs.	39,917	39,917	39,917	39,917	39,917	39,917
Mean Dependent Var.	0.21	0.21	0.21	0.21	0.21	0.21
Workplace Size	✓	✓	✓	✓	✓	✓
Fixed Effects	O+Y+E	O+Y+E+C	E+C+O*Y	O+Y+E	O+Y+E+C	E+C+O*Y

NOTE.— Top Panel: The dependent variable is the share of new hires that are female and aged 35 and below. Middle Panel: The dependent variable is the share of new hires that are female and aged 36 up to and including 50. Bottom Panel: The dependent variable is the share of new hires that are female and aged above 50.

Each observation is an occupation hired at the Employer x County-level (proxy for workplace), sampled over the years from 2006 to 2016. Columns 4-6 measure the share of ads with digital skill requirements whilst simultaneously disregarding the most common digital skill requirement for each occupation and year. Firm size is divided into buckets 1-9, 10-19, 20-49, 50-249 or 250+ employees and is measured in the previous period. Workplace x Occupation size is measured in the previous period. Standard errors are clustered at the Employer x County level. Singleton observations removed. Missing lagged variables, such as lagged firm size, is set to 0 and controlled for.

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

Table A20. *Robustness - 3 Month Hiring Window: Female Share*

	Baseline			Most Common Tech. Rem.		
Female Share, New Hires	(1)	(2)	(3)	(4)	(5)	(6)
Digital Ad Share	0.015***	0.016***	0.014**	0.019***	0.019***	0.018**
	(0.005)	(0.005)	(0.005)	(0.006)	(0.006)	(0.007)
Ln(Workplace*Occ. Size)	-0.004	-0.003	-0.002	-0.004	-0.003	-0.002
	(0.003)	(0.003)	(0.004)	(0.003)	(0.003)	(0.004)
R^2	0.24	0.24	0.25	0.24	0.24	0.25
Obs.	24,494	24,494	24,494	24,494	24,494	24,494
Mean Dependent Var.	0.78	0.78	0.78	0.78	0.78	0.78
Workplace Size	✓	✓	✓	✓	✓	✓
Fixed Effects	O+Y+E	O+Y+E+C	E+C+O*Y	O+Y+E	O+Y+E+C	E+C+O*Y

NOTE.— Each observation is an occupation hired at the Employer x County-level (proxy for workplace), sampled over the years from 2006 to 2016. Columns 4-6 measure the share of ads with digital skill requirements whilst simultaneously disregarding the most common digital skill requirement for each occupation and year. Firm size is divided into buckets 1-9, 10-19, 20-49, 50-249 or 250+ employees and is measured in the previous period. Workplace x Occupation size is measured in the previous period. Standard errors are clustered at the Employer x County level. Singleton observations removed. Missing lagged variables, such as lagged firm size, is set to 0 and controlled for.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A21. *Robustness - 3 Month Hiring Window: Age Distribution*

	Baseline			Most Common Tech. Rem.		
Share $\leq$ 35, New Hires	(1)	(2)	(3)	(4)	(5)	(6)
Digital Ad Share	0.001	0.000	0.001	-0.006	-0.006	-0.005
	(0.010)	(0.010)	(0.010)	(0.012)	(0.012)	(0.013)
Ln(Workplace*Occ. Size)	0.014***	0.014***	0.014***	0.014***	0.014***	0.014***
	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)
R^2	0.26	0.26	0.26	0.26	0.26	0.26
Obs.	24,494	24,494	24,494	24,494	24,494	24,494
Mean Dependent Var.	0.48	0.48	0.48	0.48	0.48	0.48
Workplace Size	✓	✓	✓	✓	✓	✓
Fixed Effects	O+Y+E	O+Y+E+C	E+C+O*Y	O+Y+E	O+Y+E+C	E+C+O*Y
Share 36 – 50, New Hires	(1)	(2)	(3)	(4)	(5)	(6)
Digital Ad Share	-0.000	0.000	-0.001	0.007	0.008	0.007
	(0.010)	(0.010)	(0.011)	(0.012)	(0.012)	(0.013)
Ln(Workplace*Occ. Size)	-0.006*	-0.006	-0.006	-0.006*	-0.006	-0.006
	(0.003)	(0.004)	(0.004)	(0.003)	(0.004)	(0.004)
R^2	0.18	0.18	0.19	0.18	0.18	0.19
Obs.	24,494	24,494	24,494	24,494	24,494	24,494
Mean Dependent Var.	0.35	0.35	0.35	0.35	0.35	0.35
Workplace Size	✓	✓	✓	✓	✓	✓
Fixed Effects	O+Y+E	O+Y+E+C	E+C+O*Y	O+Y+E	O+Y+E+C	E+C+O*Y
Share > 50, New Hires	(1)	(2)	(3)	(4)	(5)	(6)
Digital Ad Share	-0.001	-0.001	-0.001	-0.001	-0.002	-0.002
	(0.010)	(0.010)	(0.011)	(0.012)	(0.012)	(0.012)
Ln(Workplace*Occ. Size)	-0.008***	-0.008***	-0.008***	-0.008***	-0.008***	-0.008***
	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)
R^2	0.19	0.19	0.20	0.19	0.19	0.20
Obs.	24,494	24,494	24,494	24,494	24,494	24,494
Mean Dependent Var.	0.17	0.17	0.17	0.17	0.17	0.17
Workplace Size	✓	✓	✓	✓	✓	✓
Fixed Effects	O+Y+E	O+Y+E+C	E+C+O*Y	O+Y+E	O+Y+E+C	E+C+O*Y

NOTE.— Top Panel: The dependent variable is the share of new hires aged 35 or below. Middle Panel: The dependent variable is the share of new hires aged 36 up to and including 50. Bottom Panel: The dependent variable is the share of workers aged above 50.

Each observation is an occupation hired at the Employer x County-level (proxy for workplace), sampled over the years from 2006 to 2016. Columns 4-6 measure the share of ads with digital skill requirements whilst simultaneously disregarding the most common digital skill requirement for each occupation and year. Firm size is divided into buckets 1-9, 10-19, 20-49, 50-249 or 250+ employees and is measured in the previous period. Workplace x Occupation size is measured in the previous period. Standard errors are clustered at the Employer x County level. Singleton observations removed. Missing lagged variables, such as lagged firm size, is set to 0 and controlled for.

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

Table A22. *Robustness - 3 Month Hiring Window: Female Age Distribution*

	Baseline			Most Common Tech. Rem.		
Female Share ≤ 35, **New Hires**	(1)	(2)	(3)	(4)	(5)	(6)
Digital Ad Share	-0.005	-0.006	-0.004	-0.007	-0.007	-0.006
	(0.011)	(0.011)	(0.012)	(0.012)	(0.012)	(0.012)
Ln(Workplace*Occ. Size)	0.016***	0.016***	0.015***	0.016***	0.015***	0.015***
	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)
R^2	0.25	0.25	0.26	0.25	0.25	0.26
Obs.	22,119	22,119	22,119	22,119	22,119	22,119
Mean Dependent Var.	0.48	0.48	0.48	0.48	0.48	0.48
Workplace Size	✓	✓	✓	✓	✓	✓
Fixed Effects	O+Y+E	O+Y+E+C	E+C+O*Y	O+Y+E	O+Y+E+C	E+C+O*Y
Female Share $36-50$, **New Hires**	(1)	(2)	(3)	(4)	(5)	(6)
Digital Ad Share	0.009	0.010	0.009	0.012	0.013	0.013
	(0.011)	(0.011)	(0.012)	(0.013)	(0.013)	(0.014)
Ln(Workplace*Occ. Size)	-0.008**	-0.008**	-0.008**	-0.008**	-0.008**	-0.008**
	(0.003)	(0.004)	(0.003)	(0.003)	(0.004)	(0.003)
R^2	0.18	0.18	0.19	0.18	0.18	0.19
Obs.	22,119	22,119	22,119	22,119	22,119	22,119
Mean Dependent Var.	0.36	0.36	0.36	0.36	0.36	0.36
Workplace Size	✓	✓	✓	✓	✓	✓
Fixed Effects	O+Y+E	O+Y+E+C	E+C+O*Y	O+Y+E	O+Y+E+C	E+C+O*Y
Female Share > 50, **New Hires**	(1)	(2)	(3)	(4)	(5)	(6)
Digital Ad Share	-0.004	-0.004	-0.005	-0.006	-0.006	-0.007
	(0.012)	(0.012)	(0.013)	(0.014)	(0.014)	(0.015)
Ln(Workplace*Occ. Size)	-0.008**	-0.008***	-0.008**	-0.008**	-0.008***	-0.008**
	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)
R^2	0.19	0.19	0.20	0.19	0.19	0.20
Obs.	22,119	22,119	22,119	22,119	22,119	22,119
Mean Dependent Var.	0.16	0.16	0.16	0.16	0.16	0.16
Workplace Size	✓	✓	✓	✓	✓	✓
Fixed Effects	O+Y+E	O+Y+E+C	E+C+O*Y	O+Y+E	O+Y+E+C	E+C+O*Y

NOTE.— Top Panel: The dependent variable is the share of new hires that are female and aged 35 and below. Middle Panel: The dependent variable is the share of new hires that are female and aged 36 up to and including 50. Bottom Panel: The dependent variable is the share of new hires that are female and aged above 50.

Each observation is an occupation hired at the Employer x County-level (proxy for workplace), sampled over the years from 2006 to 2016. Columns 4-6 measure the share of ads with digital skill requirements whilst simultaneously disregarding the most common digital skill requirement for each occupation and year. Firm size is divided into buckets 1-9, 10-19, 20-49, 50-249 or 250+ employees and is measured in the previous period. Workplace x Occupation size is measured in the previous period. Standard errors are clustered at the Employer x County level. Singleton observations removed. Missing lagged variables, such as lagged firm size, is set to 0 and controlled for.

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

Table A23. *Robustness - Employer x Municipality x Occupation Matching: Female Share*

	Baseline			Most Common Tech. Rem.		
Female Share, New Hires	(1)	(2)	(3)	(4)	(5)	(6)
Digital Ad Share	0.014**	0.015**	0.013**	0.017***	0.019***	0.018**
	(0.006)	(0.006)	(0.006)	(0.007)	(0.007)	(0.007)
Ln(Workplace*Occ. Size)	-0.008***	-0.006**	-0.005	-0.008***	-0.006**	-0.004
	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)
R^2	0.24	0.25	0.26	0.24	0.25	0.26
Obs.	29,037	29,037	29,037	29,037	29,037	29,037
Mean Dependent Var.	0.78	0.78	0.78	0.78	0.78	0.78
Workplace Size	✓	✓	✓	✓	✓	✓
Fixed Effects	O+Y+E	O+Y+E+M	E+M+O*Y	O+Y+E	O+Y+E+M	E+M+O*Y

NOTE.— Each observation is an occupation hired at the Employer x Municipality-level (proxy for workplace), sampled over the years from 2006 to 2016. Columns 4-6 measure the share of ads with digital skill requirements whilst simultaneously disregarding the most common digital skill requirement for each occupation and year. Firm size is divided into buckets 1-9, 10-19, 20-49, 50-249 or 250+ employees and is measured in the previous period. Workplace x Occupation size is measured in the previous period. Standard errors are clustered at the Employer x Municipality level. Singleton observations removed. Missing lagged variables, such as lagged firm size, is set to 0 and controlled for. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A24. *Robustness - Employer x Municipality x Occupation Matching: Age Distribution*

	Baseline			Most Common Tech. Rem.		
Share ≤ 35, New Hires	(1)	(2)	(3)	(4)	(5)	(6)
Digital Ad Share	-0.003	-0.004	-0.002	-0.014	-0.015	-0.013
	(0.008)	(0.008)	(0.008)	(0.010)	(0.010)	(0.009)
Ln(Workplace*Occ. Size)	0.024***	0.022***	0.022***	0.024***	0.022***	0.021***
	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)
R^2	0.25	0.26	0.27	0.25	0.26	0.27
Obs.	29,037	29,037	29,037	29,037	29,037	29,037
Mean Dependent Var.	0.47	0.47	0.47	0.47	0.47	0.47
Workplace Size	✓	✓	✓	✓	✓	✓
Fixed Effects	O+Y+E	O+Y+E+M	E+M+O*Y	O+Y+E	O+Y+E+M	E+M+O*Y
Share $36-50$, New Hires	(1)	(2)	(3)	(4)	(5)	(6)
Digital Ad Share	0.006	0.007	0.004	0.015**	0.015**	0.013*
	(0.007)	(0.007)	(0.007)	(0.007)	(0.007)	(0.007)
Ln(Workplace*Occ. Size)	-0.017***	-0.016***	-0.016***	-0.017***	-0.015***	-0.015***
	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)
R^2	0.17	0.19	0.19	0.17	0.19	0.19
Obs.	29,037	29,037	29,037	29,037	29,037	29,037
Mean Dependent Var.	0.36	0.36	0.36	0.36	0.36	0.36
Workplace Size	✓	✓	✓	✓	✓	✓
Fixed Effects	O+Y+E	O+Y+E+M	E+M+O*Y	O+Y+E	O+Y+E+M	E+M+O*Y
Share > 50, New Hires	(1)	(2)	(3)	(4)	(5)	(6)
Digital Ad Share	-0.002	-0.002	-0.002	-0.001	-0.000	-0.000
	(0.008)	(0.008)	(0.008)	(0.008)	(0.008)	(0.007)
Ln(Workplace*Occ. Size)	-0.007***	-0.006**	-0.006**	-0.007***	-0.006**	-0.006**
	(0.003)	(0.003)	(0.002)	(0.003)	(0.003)	(0.002)
R^2	0.18	0.20	0.20	0.18	0.20	0.20
Obs.	29,037	29,037	29,037	29,037	29,037	29,037
Mean Dependent Var.	0.17	0.17	0.17	0.17	0.17	0.17
Workplace Size	✓	✓	✓	✓	✓	✓
Fixed Effects	O+Y+E	O+Y+E+M	E+M+O*Y	O+Y+E	O+Y+E+M	E+M+O*Y

NOTE.— Top Panel: The dependent variable is the share of new hires aged 35 or below. Middle Panel: The dependent variable is the share of new hires aged 36 up to and including 50. Bottom Panel: The dependent variable is the share of workers aged above 50. Each observation is an occupation hired at the Employer x Municipality-level (proxy for workplace), sampled over the years from 2006 to 2016. Columns 4-6 measure the share of ads with digital skill requirements whilst simultaneously disregarding the most common digital skill requirement for each occupation and year. Firm size is divided into buckets 1-9, 10-19, 20-49, 50-249 or 250+ employees and is measured in the previous period. Workplace x Occupation size is measured in the previous period. Standard errors are clustered at the Employer x Municipality level. Singleton observations removed. Missing lagged variables, such as lagged firm size, is set to 0 and controlled for.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A25. *Robustness - Employer x Municipality x Occupation Matching: Female Age Distribution*

	Baseline			Most Common Tech. Rem.		
Female Share ≤ 35, **New Hires**	(1)	(2)	(3)	(4)	(5)	(6)
Digital Ad Share	-0.009	-0.011	-0.008	-0.015	-0.016*	-0.013
	(0.008)	(0.009)	(0.009)	(0.010)	(0.009)	(0.009)
Ln(Workplace*Occ. Size)	0.025***	0.022***	0.022***	0.025***	0.022***	0.022***
	(0.005)	(0.004)	(0.004)	(0.005)	(0.004)	(0.004)
R^2	0.24	0.26	0.27	0.24	0.26	0.27
Obs.	26,336	26,336	26,336	26,336	26,336	26,336
Mean Dependent Var.	0.47	0.47	0.47	0.47	0.47	0.47
Workplace Size	✓	✓	✓	✓	✓	✓
Fixed Effects	O+Y+E	O+Y+E+M	E+M+O*Y	O+Y+E	O+Y+E+M	E+M+O*Y
Female Share $36-50$, **New Hires**	(1)	(2)	(3)	(4)	(5)	(6)
Digital Ad Share	0.014*	0.015*	0.013	0.018**	0.019**	0.017**
	(0.008)	(0.008)	(0.008)	(0.008)	(0.008)	(0.008)
Ln(Workplace*Occ. Size)	-0.018***	-0.016***	-0.016***	-0.018***	-0.016***	-0.016***
	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)
R^2	0.17	0.19	0.19	0.17	0.19	0.19
Obs.	26,336	26,336	26,336	26,336	26,336	26,336
Mean Dependent Var.	0.37	0.37	0.37	0.37	0.37	0.37
Workplace Size	✓	✓	✓	✓	✓	✓
Fixed Effects	O+Y+E	O+Y+E+M	E+M+O*Y	O+Y+E	O+Y+E+M	E+M+O*Y
Female Share > 50, **New Hires**	(1)	(2)	(3)	(4)	(5)	(6)
Digital Ad Share	-0.004	-0.004	-0.005	-0.004	-0.003	-0.004
	(0.008)	(0.008)	(0.008)	(0.009)	(0.009)	(0.009)
Ln(Workplace*Occ. Size)	-0.007**	-0.006**	-0.006**	-0.007**	-0.006**	-0.006**
	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)
R^2	0.18	0.20	0.20	0.18	0.20	0.20
Obs.	26,336	26,336	26,336	26,336	26,336	26,336
Mean Dependent Var.	0.16	0.16	0.16	0.16	0.16	0.16
Workplace Size	✓	✓	✓	✓	✓	✓
Fixed Effects	O+Y+E	O+Y+E+M	E+M+O*Y	O+Y+E	O+Y+E+M	E+M+O*Y

NOTE.— Top Panel: The dependent variable is the share of new hires that are female and aged 35 and below. Middle Panel: The dependent variable is the share of new hires that are female and aged 36 up to and including 50. Bottom Panel: The dependent variable is the share of new hires that are female and aged above 50. Each observation is an occupation hired at the Employer x Municipality-level (proxy for workplace), sampled over the years from 2006 to 2016. Columns 4-6 measure the share of ads with digital skill requirements whilst simultaneously disregarding the most common digital skill requirement for each occupation and year. Firm size is divided into buckets 1-9, 10-19, 20-49, 50-249 or 250+ employees and is measured in the previous period. Workplace x Occupation size is measured in the previous period. Standard errors are clustered at the Employer x Municipality level. Singleton observations removed. Missing lagged variables, such as lagged firm size, is set to 0 and controlled for. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A26. *Robustness - Employer x Occupation Matching: Female Share*

	Baseline		Most Common Tech. Rem.	
Female Share, New Hires	(1)	(2)	(3)	(4)
Digital Ad Share	0.007	0.004	0.011	0.008
	(0.007)	(0.007)	(0.007)	(0.007)
Ln(Employer*Occ. Size)	-0.002	-0.000	-0.001	-0.000
	(0.003)	(0.003)	(0.003)	(0.003)
R^2	0.26	0.27	0.26	0.27
Obs.	27,255	27,255	27,255	27,255
Mean Dependent Var.	0.78	0.78	0.78	0.78
Employer Size	✓	✓	✓	✓
Fixed Effects	O+Y+E	E+O*Y	O+Y+E	E+O*Y

NOTE.— Each observation is an occupation hired at the employer-level, sampled over the years from 2006 to 2016. Columns 4-6 measure the share of ads with digital skill requirements whilst simultaneously disregarding the most common digital skill requirement for each occupation and year. Firm size is divided into buckets 1-9, 10-19, 20-49, 50-249 or 250+ employees and is measured in the previous period. Employer x Occupation size is measured in the previous period. Standard errors are clustered at the Employer level. Singleton observations removed. Missing lagged variables, such as lagged firm size, is set to 0 and controlled for. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A27. *Robustness - Employer x Occupation Matching: Age Distribution*

	Baseline		Most Common Tech. Rem.	
Share ≤ 35, New Hires	(1)	(2)	(3)	(4)
Digital Ad Share	-0.006	-0.004	-0.013	-0.010
	(0.009)	(0.008)	(0.009)	(0.009)
Ln(Employer*Occ. Size)	0.016***	0.016***	0.016***	0.015***
	(0.003)	(0.003)	(0.003)	(0.003)
R^2	0.28	0.28	0.28	0.28
Obs.	27,255	27,255	27,255	27,255
Mean Dependent Var.	0.47	0.47	0.47	0.47
Employer Size	✓	✓	✓	✓
Fixed Effects	O+Y+E	E+O*Y	O+Y+E	E+O*Y
Share $36-50$, New Hires	(1)	(2)	(3)	(4)
Digital Ad Share	0.007	0.005	0.014	0.011
	(0.008)	(0.008)	(0.009)	(0.009)
Ln(Employer*Occ. Size)	-0.009***	-0.009***	-0.009***	-0.008***
	(0.003)	(0.003)	(0.003)	(0.003)
R^2	0.20	0.21	0.20	0.21
Obs.	27,255	27,255	27,255	27,255
Mean Dependent Var.	0.36	0.36	0.36	0.36
Employer Size	✓	✓	✓	✓
Fixed Effects	O+Y+E	E+O*Y	O+Y+E	E+O*Y
Share > 50, New Hires	(1)	(2)	(3)	(4)
Digital Ad Share	-0.000	-0.001	-0.001	-0.001
	(0.006)	(0.006)	(0.007)	(0.007)
Ln(Employer*Occ. Size)	-0.007**	-0.007***	-0.007**	-0.007***
	(0.002)	(0.002)	(0.002)	(0.002)
R^2	0.21	0.22	0.21	0.22
Obs.	27,255	27,255	27,255	27,255
Mean Dependent Var.	0.17	0.17	0.17	0.17
Employer Size	✓	✓	✓	✓
Fixed Effects	O+Y+E	E+O*Y	O+Y+E	E+O*Y

NOTE.— Top Panel: The dependent variable is the share of new hires aged 35 or below. Middle Panel: The dependent variable is the share of new hires aged 36 up to and including 50. Bottom Panel: The dependent variable is the share of workers aged above 50. Each observation is an occupation hired at the employer-level, sampled over the years from 2006 to 2016. Columns 4-6 measure the share of ads with digital skill requirements whilst simultaneously disregarding the most common digital skill requirement for each occupation and year. Firm size is divided into buckets 1-9, 10-19, 20-49, 50-249 or 250+ employees and is measured in the previous period. Employer x Occupation size is measured in the previous period. Standard errors are clustered at the Employer level. Singleton observations removed. Missing lagged variables, such as lagged firm size, is set to 0 and controlled for. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A28. *Robustness - Employer x Occupation Matching: Female Age Distribution*

	Baseline		Most Common Tech. Rem.	
Female Share $\leq$ 35, New Hires	(1)	(2)	(3)	(4)
Digital Ad Share	-0.010	-0.008	-0.014	-0.012
	(0.009)	(0.009)	(0.010)	(0.010)
Ln(Workplace*Occ. Size)	0.017***	0.017***	0.017***	0.017***
	(0.003)	(0.003)	(0.003)	(0.003)
R^2	0.27	0.27	0.27	0.27
Obs.	24,806	24,806	24,806	24,806
Mean Dependent Var.	0.47	0.47	0.47	0.47
Workplace Size	✓	✓	✓	✓
Fixed Effects	O+Y+E	E+O*Y	O+Y+E	E+O*Y
Female Share $36-50$**, New Hires**	(1)	(2)	(3)	(4)
Digital Ad Share	0.014	0.013	0.020**	0.018*
	(0.009)	(0.009)	(0.010)	(0.010)
Ln(Workplace*Occ. Size)	-0.011***	-0.011***	-0.011***	-0.011***
	(0.003)	(0.003)	(0.003)	(0.003)
R^2	0.20	0.20	0.20	0.20
Obs.	24,806	24,806	24,806	24,806
Mean Dependent Var.	0.37	0.37	0.37	0.37
Workplace Size	✓	✓	✓	✓
Fixed Effects	O+Y+E	E+O*Y	O+Y+E	E+O*Y
Female Share $>$ 50**, New Hires**	(1)	(2)	(3)	(4)
Digital Ad Share	-0.004	-0.004	-0.006	-0.006
	(0.007)	(0.007)	(0.007)	(0.007)
Ln(Workplace*Occ. Size)	-0.006**	-0.006**	-0.006**	-0.006**
	(0.002)	(0.002)	(0.002)	(0.002)
R^2	0.21	0.21	0.21	0.21
Obs.	24,806	24,806	24,806	24,806
Mean Dependent Var.	0.17	0.17	0.17	0.17
Workplace Size	✓	✓	✓	✓
Fixed Effects	O+Y+E	E+O*Y	O+Y+E	E+O*Y

NOTE.— Top Panel: The dependent variable is the share of new hires that are female and aged 35 and below. Middle Panel: The dependent variable is the share of new hires that are female and aged 36 up to and including 50. Bottom Panel: The dependent variable is the share of new hires that are female and aged above 50. Each observation is an occupation hired at the employer-level, sampled over the years from 2006 to 2016. Columns 4-6 measure the share of ads with digital skill requirements whilst simultaneously disregarding the most common digital skill requirement for each occupation and year. Firm size is divided into buckets 1-9, 10-19, 20-49, 50-249 or 250+ employees and is measured in the previous period. Employer x Occupation size is measured in the previous period. Standard errors are clustered at the Employer level. Singleton observations removed. Missing lagged variables, such as lagged firm size, is set to 0 and controlled for. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Essay II. Software, Robots, AI & The Gender Wage Gap: Evidence From US Cities

Acknowledgements: I would like to express my gratitude to Oskar Nordström Skans, Arizo Karimi and Josefin Videnord, as well as the seminar participants at the Uppsala Center for Labor Studies, SUDSWEC 2022, EALE 2023, and the Labor Group at Uppsala University for their helpful comments and suggestions. I gratefully acknowledge funding from WASP-HS.

1 Introduction

The gender wage gap in all OECD countries has seen a sharp decline from the 1980s and onward. In the US, the unadjusted gender wage gap has halved between 1980 and 2010 (Blau and Kahn, 2017). Both supply- and demand-side factors in the labor market contribute to this decline. In this paper, I focus on the role of technological advancements, one of the most important trends affecting labor demand in the past decades.

New technologies replace tasks and occupations, which may affect male and female workers differently due to occupational gender segregation (Autor et al., 2003; Lindley, 2012). I use a direct measurement of technology exposure at the occupation level, based on task-replacing patents, and analyze how resulting changes in labor demand have affected the gender wage gap. To this end, I leverage cross-city differences in employment shares of robot and software exposed occupations and analyze how technology-induced changes in labor demand have affected the evolution of the gender wage gap in different cities. By exploiting the cross-city differences in initial employment shares of highly exposed occupations, prior to large scale software or robot adoption, I can establish a relationship between city-level technology exposure and development of the gender wage gap from 1980 to 2010. Extrapolating from these historic data on the effects of robot and software technology, I also comment on how exposure to Artificial Intelligence (AI) may affect the future development of the gender wage gap in the following decades.

Throughout the study period (1980 - 2010), gender-related labor market disparities have decreased significantly, and the literature studying this trend is very large (Goldin, 2006). Increase in female educational attainment and other supply-side factors have played crucial roles in the overall decline in the gender wage gap. [1]. However, research also points to an important role for demand side factors. During the 1980s and 1990s, demand shifts favored female-dominated occupations (such as clerical work) and social and interpersonal skills, predominantly associated with female workers (Blau and Kahn, 2013; Cortes et al., 2018). The decline in the manufacturing sector during the late 1970s and 1980s mainly affected male workers (Blau and Kahn, 1997). Part of the decline in manufacturing is attributed to the large-scale introduction of industrial robots that automate tasks and replace labor, thereby decreasing labor demand (Blau and Kahn, 2006). Workplace computerization during the 1990s has also been attributed to decreasing the gender wage gap through complementarity with female-dominated tasks and skills (Beaudry and Lewis, 2014; Spitz-Oener, 2006).

Thus, male and female workers are not equally affected by technological advancements due to occupational gender segregation (Blau and Kahn, 2017;

[1] For references on family-friendly policies see Blau and Kahn (2013). For anti-discriminatory legislation Holzer and Neumark (2000). For education, see (Becker et al., 2010; Goldin et al., 2006). For structural changes, see Blau and Kahn (2006) and Blau et al. (2013)

Blau and Kahn, 2013). The effect of technology adoption on employment rates and average wages also differs between local labor markets, due to the regional clustering of occupations and industries. Acemoglu and Restrepo (2020) find that employment and relative wages have fallen in US cities with (initially) large employment shares in industries that were prone to robot adoption.

Only a few studies have explicitly linked technology adoption to the gender wage gap. Some important examples include Cortes et al. (2020), Beaudry and Lewis (2014), and Black and Spitz-Oener (2010), which are all closely related to this paper. Cortes et al. (2020) use data from the US and Portugal (1993-2019) and attribute the decrease in the gender wage gap to the overrepresentation of men in declining routine manual occupations, but without using direct measures of exposure to technology. Instead, technology exposure is assumed through the intensity of routine tasks in occupations, previously shown to be primarily affected by technology-induced labor displacement. I use a direct measurement of technology exposure that allows me to distinguish between different types of technology, such as robot or software technology, that may affect male and female workers differently. Black and Spitz-Oener (2010) use data from West Germany (1979-1999) to show that a large share of the decrease in the gender wage gap can be attributed to the increase in the relative price of nonroutine and cognitive tasks, mainly populated by the female labor force, induced by the increased computerization of workplaces. Black and Spitz-Oener (2010) use a direct measurement of technology exposure (computer adoption in the workplace), but do not study the direct impact on the gender wage gap. Instead, they analyze the gender-specific effects of changes in skill and task prices. Changes in the gender wage gap capture both the change in task prices and changes in employment shares, both accelerated by technological change. Beaudry and Lewis (2014) use US data (1980-2010) and initial cross-city differences in college attainment to explain differences in the decrease in the gender wage gap between cities, attributed to the increased college premium expedited by workplace computerization. Similar to my approach, they exploit initial differences in labor force characteristics between cities to capture propensity differences in technology-induced labor demand changes.

The role of technology adoption on the gender wage gap is addressed using patent-based measurements of technology exposure at the occupation level drawn from Webb (2019). These data are linked to information about the occupational structure at the city level and the evolution of the gender wage gap from the American Community Survey (ACS) (Ruggles et al., 2021). My approach follows Acemoglu and Restrepo (2020) who adopt a cross-region analysis, using US commuting zones, and find that commuting zones with larger employment shares in industries with high rates of robot adoption see declines in the employment-to-population ratio and average wages. There are two key differences between Acemoglu and Restrepo (2020) and this paper: I focus on robot and software exposure across occupations instead of industry-

level robot adoption. Furthermore, I focus on the gender wage gap, not wage levels.

This study produced several finding. I begin by documenting that male workers are overrepresented in software and robot exposed occupations. Furthermore, labor demand in these occupations declined from 1980 to 2010, resulting in lower average wage growth and declining employment shares. I show that cities with initially (1980) higher employment shares of software and robot exposed occupations experienced a greater decline in the gender wage gap between 1980 and 2010. Thus, labor markets that were affected by technology adoption became more gender equal through technology-induced labor displacement affecting male workers more than female workers. Software technology has had the most distinct effect on the gender wage gap, most likely due to on average higher wages in software-exposed occupations compared to robot-exposed occupations. The results are robust to controlling for city-level demographics and to the inclusion of controls for changes in the gender wage gap prior to large scale adoption of software and robot technology.

I further analyze the relationship between AI-patents and occupations in an attempt to assess how AI will affect the gender wage gap in the coming decades. AI and software patenting activity are positively correlated at the occupational level but AI-exposed occupations tend to be higher up in the wage distribution. Similar to robot- and software-exposed occupations, AI-occupations are male-dominated. At the city level, the correlation between software exposure in 1980 and AI exposure in 2019 is very strong. Assuming similar labor-replacement effects of AI technology, this suggests that the gender wage gap in these cities could continue to decline.

Overall, the results show that the adoption of robot and software-related technologies has played a significant role in the development of the gender wage gap between 1980 and 2010. Due to occupational gender segregation, male and female workers do not experience the labor market consequences of technology adoption equally, resulting in crucial shifts in the gender wage gap. My results further indicate that labor markets with many workers in occupations with high rates of AI technology exposure may see larger declines in the gender wage gap in the future.

The rest of the paper is organized as follows. Section 2 presents a simple theoretical introduction to how technology exposure affects the gender wage gap through changes in labor demand. Section 3 introduces data sources and measurements. Section 4 presents US-level demographics of workers most exposed to software and robot technology and changes in labor demand from 1980 to 2010. Section 5 presents my main empirical results from the city-level analysis, where I regress city-level changes in the gender wage gap on employment shares of the most exposed occupations, for both software and robot technology. Section 6 presents a tentative analysis of AI exposure and the gender wage gap. Section 7 concludes. The appendix includes robustness

checks and additional information on the demographics of the most exposed workers.

2 Technology and The Gender Wage Gap

To fix ideas, this section describes the potential mechanisms through which technology-induced changes in labor demand affects the gender wage gap, depending on the relative wages of exposed occupations and the rate of occupational gender segregation.

The gender wage gap in city c, in year t, is defined as

$$\begin{aligned}\text{Gender Wage Gap}_{ct} &= ln(\overline{\text{Wage}_{ct}^{Male}}) - ln(\overline{\text{Wage}_{ct}^{Female}}) \\ &= ln(\sum_j S_{ctj}^{Male} * \overline{Wage_{ctj}^{Male}}) - ln(\sum_j S_{ctj}^{Female} * \overline{Wage_{ctj}^{Female}})\end{aligned}$$

Where S_{ctj}^{Male} and S_{ctj}^{Female} is the share of the male and female labor force in occupation j, in city or country c and in year t. $\overline{Wage_{ctj}^{Male}}$ and $\overline{Wage_{ctj}^{Female}}$ denote the average wage in occupation j for male and female workers, respectively. Through automation or productivity enhancements, technology adoption has the potential to affect both the employment shares and average wages of occupations, and thereby both components of the gender wage gap. Declines in demand for occupations, driven by technological automation of tasks, decrease the share of workers in that occupation and the average wages of the workers who remain. [2] Conversely, labor-augmenting innovations will increase the demand for affected occupations, increasing both employment shares and wages. [3]

In this paper, technology exposure is measured by the overlap of technological capabilities, as defined by patent descriptions, and occupational tasks, as described by O*NET, and therefore, technology exposure is here seen for its ability to replace labor. High levels of technology exposure, as measured in this paper, are therefore likely to incur declines in labor demand, resulting in declines in average wages and a resorting of workers into other occupations or into unemployment.

Due to significant occupational gender segregation, any changes in labor demand as a result of technology adoption will affect male and female workers

[2] Previous work has primarily linked this to the routineness of tasks and their suitability for automation (Autor et al., 2003; Goos and Manning, 2007; Acemoglu and Autor, 2011; Cortes, 2016).

[3] The introduction of personal computers and the compatibility with high-skill tasks has previously been raised as a source of the increasing college premium and overall wage inequality (Acemoglu, 2003; Beaudry et al., 2010). Acemoglu and Restrepo (2018) build a framework through which new technologies increase labor demand.

differently and therefore affect the gender wage gap. [4] Technology-induced declines in labor demand will affect equilibrium wages through the following three mechanisms:

1. **Direct effect** - Declines in average wages for workers who remain in the exposed occupations
2. **Sorting effect** - When exposed occupations pay above average wages, a reallocation of labor to occupations with lower wages
3. **Indirect supply effect** - Indirect effects on workers in "secondary" occupations, into which displaced workers allocate. For affected workers in occupations with low average wages, decreased labor demand moves workers to other low-wage occupations. This reduces the average wages among both incumbent and new workers in occupations less exposed to technology.

If the most technology-exposed occupations are male-dominated, declines in the labor demand for these occupations will decrease the gender wage gap through one or more of the three channels mentioned above. Which ones and to what degree will depend on where in the wage distribution the exposed occupations occur for different types of technologies. Firstly, automation of tasks decreases the average wage for workers who still remain in the occupation. All else being equal, this lowers the average male wage if the share of male labor force employed in exposed occupations is larger than the share of the female labor force. This mechanism plays out regardless of where in the wage distribution the exposed occupations are. Secondly, if labor demand decreases for male-dominated mid- to high-wage occupations, this pushes workers into occupations with on average lower wages. This mechanism will materialize only for software-exposed occupations, and to an even greater extent, AI-exposed occupations, that occur primarily in the middle and higher up in the wage distribution, respectively. Lastly, if labor demand decreases in male-dominated low-wage occupations, as is the case with robot-exposed occupations, this pushes workers to other low-wage occupations, increasing labor supply and potentially lowering average wages, for both incumbent and new workers. All three channels, but to different extents depending on the type of technology, drive down the gender wage gap by decreasing the average wage in male-dominated occupations and moving large groups of male workers into occupations with comparatively lower wages. Female workers in these occupations also stand to be affected by the declines in labor demand, but they make up a much smaller share of the affected labor force, which ultimately nets a decline in the gender wage gap. Lastly, declines in labor demand can also force workers out of the labor market all together. This can potentially change the demographic make-up of the labor force, both in terms of gender and, for example, educational background.

[4] For references on occupational gender segregation, see Blau et al. (2013), Blackburn et al. (2000), and Hegewisch and Hartmann (2014)

Due to industrial and occupational clustering over cities and local labor markets, technological advancements will affect different labor markets to different extents. Similarly to the approaches of Acemoglu and Restrepo (2020) and Autor and Dorn (2013), I exploit these differences in technology exposure rates at the local labor market level. Technology exposure, as measured in this context, occurs at the occupation level, but the gender wage gap is ultimately an equilibrium outcome. Local labor markets with high relative employment shares of technology exposed occupations should see larger effects of declining labor demand, affecting mainly male workers, and thus the gender wage gap should decline more in these labor markets. A geographically defined local labor market, such as a metropolitan area, holds the pool of workers, both in and out of employment, similar over time as opposed to analyzing the worker composition of industries, for example.

3 Data Sources and Measurements

In this section, I describe the data sources used in this paper, describe the measurement of technology exposure, and present summary statistics of the final sample. I use a patent-based technology exposure measurement created by Webb (2019) that measures the overlap in tasks performed by patented technology and tasks performed at the occupation level. The patents are collected from Google Patents Public Data and contain the universe of patents filed from 17 major patenting offices up until 2018. I combine this with data on occupational employment shares and average wages from the American Community Survey from 1980 to 2010 in my main sample for the robot and software analysis. For the AI exposure analysis, I use employment and wage data from 2019. At a later stage in the analysis, I aggregate my data to city level, using metropolitan areas as a proxy for local labor markets. I complement the main sample with data from 1950 and 1970 to analyze pretrends in employment shares, wages, and city-level gender wage gaps.

Exposure Measurement

Three types of technologies are considered in this paper; robot, software, and AI. To measure technology exposure, I use patent-based exposure scores, created by Webb (2019), which measure the extent to which patents filed can perform tasks within each occupation. [5] Patent data is collected from Google Patents Public Data provided by IFI CLAIMS Patent Services. This data contains all patents from 17 major patenting offices, including the United States Patent and Trademark Office (USPTO) and the European Patent Office (EPO). The database is updated quarterly and stretches back to the late 18th century

[5] Similar approaches using patent data to calculate exposure to technology are found in Dechezleprêtre et al. (2021), Kogan et al. (2021), and Mann and Püttmann (2023) and most recently, Autor et al. (2024).

for some patent offices. For the scores used in this paper, all patents up to and including 2018 are used. From the patent titles, Webb (2019), extracts all verb-noun pairs, such as 'analyze data' or 'make diagnosis', and all common synonyms for these. These are then used to measure the rate of overlap between tasks performed in occupations and tasks performed by patented technology. Occupations' tasks are collected from the O*NET database provided by the US Department of Labor, which details what tasks are most commonly performed by a specific occupation. These are narrowed down to one or more verb-noun pairs along with the most common synonyms, applying the same process as with the patent titles. The O*NET also contains a measurement of how important a task is for a given occupation, relative other tasks, which are used as weights when calculating the overlap. For each technology t and verb-noun pair c in the patent data, the relative frequency, rf_c^t is defined as

$$rf_c^t = \frac{f_c^t}{\sum_{c \in C^t} f_c^t} \tag{1}$$

where C^t is the set of all verb-noun pairs extracted from patent titles for technology t. For each occupation o, the overall exposure to technology t is given by

$$Exposure_{o,t} = \frac{\sum_{k \in K_o} [\omega_{k,o} \times \sum_{c \in S_k} rf_c^t]}{\sum_{k \in K_o} [\omega_{k,o} \times |\{c : c \in S_k\}|]} \tag{2}$$

where K_o is the set of tasks in occupation o, S_k is the set of verb-noun pairs extracted from task k and $\omega_{k,o}$ is the importance weight of task k. This exposure score then captures the patenting activity in technology t directed toward occupation o, weighted by the importance of the task for that specific occupation. Each occupation, using Dorn et al. (2009)'s occupational code, is then assigned a percentile score which is what I will be using throughout the analysis. I also create quartiles of exposure for each technology type, from least to most exposed, defined so that each quartile covers a quarter of all workers from 1980 in my final weighted sample. The quartile assignments for each occupation will then carry over to each time period in my analysis, regardless of how workers redistribute over time.

The most common types of tasks extracted from robot patents include the verbs 'clean', 'control', 'weld' and 'move', resulting in such tasks as 'move object' and 'weld component'. The most common types of tasks extracted from software patents include the verbs 'record', 'store', 'control' and 'reproduce', resulting in such tasks as 'record data' and 'control access'. Lastly, the most common types of tasks extracted from AI patents include the verbs 'recognize', 'predict', 'detect' and 'identify', resulting in such tasks as 'recognize patterns' and 'detect abnormality'. [6] See Table A1 in the Appendix for the most exposed and least exposed occupations in all categories of technology.

[6] For a full list of the most common verb-noun pairs for each technology type, see Webb (2019).

One important strength with this measurement is that it simultaneously captures both the technological frontier, through nonzero exposure scores, and technological availability, since the scores are higher when several different patents, within the same technology category, all claim to be able to perform a given task.

There are also some potential limitations in using this specific measurement. First, exposure measurements are only recorded in one time period, 2018, and do not vary over time. I can therefore not capture the effect of *changes* in exposure over time and use instead the measurement as potential exposure, especially when looking at wages and employment in my initial period, 1980, where the exposure measurement is viewed as future patenting activity that may affect these occupations. Secondly, the measurement does not capture the interactive nature of, for example, industry and occupation, and instead assumes that all occupations are equally affected regardless of industry and geographical location. I address this by limiting my sample to only metropolitan areas to avoid rural vs. urban comparisons and assume similar technology adoption over cities, given their metropolitan status.

By using this definition of technology exposure, I follow the task-based approach, as defined by Autor et al. (2003) and further developed by Acemoglu and Restrepo (2018), where technology replaces human labor in tasks within occupations. Although technology implementation can act as both a substitute and a complement to human labor, this measurement is assumed to measure replacement risk and/or automation risk. However, it is not certain that this measurement always needs to be negative in terms of wages or employment for occupations with high levels of exposure since automation in some tasks frees up more time for other tasks that may be more lucrative/productive or that the technology is used as a tool to increase the productivity of tasks without replacing human labor.

Wage and Employment Data

I use individual-level microdata from the US Census and the American Community Survey (ACS) from 1980 to 2019 provided by IPUMS (Ruggles et al., 2021). The ACS replaced the long-form US Census survey that was sent out to 1-in-6 households every 10 years, and is now, since 2001, a yearly survey sent out to a smaller sample of households. The US Census and the ACS are repeated cross-sectional surveys sent out to American households that contain questions about work, income, and educational background, among other things. The 1980, 1990 and 2000 samples are 1-in-20 national random samples while the 2010 and 2019 samples are 1-in-100 national random samples. Combined with the survey person weights provided by the US Census and the ACS, the data provide nationally and locally representative samples. The time span of 1980 to 2010 is chosen following Beaudry and Lewis (2014), such that I define 1980 as the "Pre-PC" era. Although computers did exist in the market at this point, they did so in a very limited capacity. After this point, computer

usage increased rapidly, where, from 1984 to 1993, computer use among all workers increased from approximately a quarter to approximately half in less than a decade (Autor et al., 1998). For this reason, I will use 1980 as the initial time period throughout the analysis. For the main analysis, I restrict my sample to individuals between 16 to 64 years of age, who are employed and drop unpaid family workers. I further restrict my sample to individuals who live in metropolitan areas that appear in both of my main time periods, 1980 and 2010. Due to the change in sampling size from 1980 to 2010, some cities do not appear in the later samples. I use the cross-walk provided by Dorn et al. (2009) to translate the ACS occupational codes to the occupational codes used in the technology exposure data. To calculate city*occupation-level employment and average wages I use the IPUMS survey person weight adjusted to full-time full-year equivalence. Missing labor supply weights are imputed based on occupation and educational background. Wages are also recalculated to full-time, full-year equivalence. As in Dorn et al. (2009), the top-coded earnings are adjusted by a factor of 1.5. Lastly, wages are winsorized at the bottom 1%.

Summary Statistics

The final sample includes 324 occupations and 239 metropolitan areas. The summary statistics are presented in Table 1. The number of unweighted observations drop from 1980 to 2010 due to the decrease in the sample size of the surveys from a 5% representative sample to a 1% representative sample of the US population. The summary statistics reveal several notable patterns. Firstly, they demonstrate the narrowing of the college graduation gap between men and women over time. There is also a decrease in the gender wage gap, defined as the logarithmic difference between the average weekly wages of men and women, from around 54% to 21%. [7] Lastly, the share of workers, both male and female, in the quartile most exposed to software technology decreases over time. In contrast, the share of workers in the least exposed quartile increases. This means that the relative employment shares of the most exposed occupations have decreased from 1980 to 2010. The same trends appear in the employment shares of occupations exposed to robot technology.

[7] Histograms of log weekly wages for men and women in 1980 and 2010 are presented in Figure A1

Table 1. *Summary statistics - Sample Means (1980 - 2010)*

	1980 Women	1980 Men	2010 Women	2010 Men
Age	36.28	37.49	40.91	40.72
Ethnicity				
Share Black	.13	.1	.15	.1
Share Hispanic	.06	.07	.15	.19
Educational Background				
Share Less Than High School	.18	.22	.06	.09
Share High School	.39	.32	.29	.32
Share Some College	.26	.24	.27	.24
Share College	.17	.23	.38	.35
Weekly Wage				
Average Ln(Weekly Wage)	6.03	6.46	6.36	6.55
P10 Ln(Weekly Wage)	5.39	5.69	5.51	5.61
Median Ln(Weekly Wage)	6.08	6.54	6.38	6.53
P90 Ln(Weekly Wage)	6.69	7.16	7.2	7.45
Software Exposure Quartiles				
Least Exposed Q1	.31	.22	.35	.25
Medium Low Exposed Q2	.37	.14	.37	.2
Medium High Exposed Q3	.19	.31	.19	.3
Most Exposed Q4	.13	.32	.08	.25
Robot Exposure Quartiles				
Least Exposed Q1	.3	.23	.43	.3
Medium Low Exposed Q2	.32	.21	.26	.21
Medium High Exposed Q3	.24	.23	.21	.22
Most Exposed Q4	.14	.33	.1	.28
Industry Shares				
Agriculture	.01	.01	.01	.02
Mining	0	.01	0	.01
Construction	.01	.08	.01	.09
Manufacturing	.18	.3	.08	.16
Transportation	.05	.11	.05	.1
Wholesale trade	.03	.06	.02	.04
Retail trade	.16	.12	.16	.16
Finance	.1	.05	.09	.06
Repair	.03	.04	.05	.08
Personal Services	.04	.01	.04	.02
Entertainment	.01	.01	.01	.02
Professional	.31	.12	.43	.18
Public Administration	.06	.07	.06	.06
Obs.	1,577,761	1,835,845	466,941	466,904
Weighted Obs.	26,014,329	38,545,602	47,840,446	55,503,852

NOTE.— Summary statistics presenting sample means for male and female workers in 1980 and 2010. Sample restricted to employed individuals, age 16-64, in metropolitan areas that exist in both the 1980 and 2010 sample. Observations are weighted using ACS person weights adjusted to full-time full-year equivalence. Wages are winsorized at the bottom 1%. 1980 wages are adjusted for inflation using 2010 prices. Exposure quartiles defined to include 1/4 of all workers in 1980 weighted sample, occupation-level quartile assignments carry over to 2010 sample. US Census & American Community Survey (1980-2010) (Ruggles et al., 2021). Exposure Measurement provided by Webb (2019).

4 Worker Demographics, Labor Demand and Technology Exposure

In this section, I show the types of workers most exposed to software and robot technology and how labor demand has changed for these workers from 1980 to 2010 compared to those in less exposed occupations. I focus on robot and software technology specifically at this stage and return to AI technology in a later part of the paper.

4.1 Worker Demographics and Technology Exposure

Gender segregation in the labor market means that men and women work, on average, in different occupations. Since exposure to software and robotic technology differs between occupations, men and women are at different risk of technology-induced labor displacement. Furthermore, the relative wage levels of the most exposed occupations will determine the direct and indirect effect that potential labor displacement will have on the gender wage gap.

Table 2 presents summary statistics for occupations in the most and least exposed software and robot exposure quartiles in 1980. The quartiles are created separately for each type of technology and defined to include 25% of all workers in the weighted sample. Occupations in the highest software or robot exposure quartile have, on average, a smaller female share of labor supply, smaller share of college educated, larger share black or hispanic workers, and occur primarily in the manufacturing industry. Conversely, occupations in the lowest software or robot exposure quartile have, on average, a larger female share, a large share of college educated workers, and occur primarily in the professional and retail trade industries. In conclusion, any adverse labor market effects of technology exposure will primarily affect male, less educated, workers in manufacturing. Robot-exposed occupations have on average lower wages compared to less robot-exposed occupations. Software-exposed occupations, on the other hand, have slightly higher average wages, compared to less software-exposed occupations.

Table 2. *Summary Statistics - Occupation-Level Means over Technology Exposure Quartiles (1980)*

	Software Exposure		Robot Exposure	
Exposure Quartile	Q1	Q4	Q1	Q4
Female Share	.48	.21	.47	.22
Avg. Age	36.6	36.63	37.44	36.48
Weekly Wage				
Avg. Ln(Weekly Wage)	6.24	6.28	6.4	6.18
P10 Ln(Weekly Wage)	5.42	5.55	5.65	5.4
Median Ln(Weekly Wage)	6.27	6.35	6.44	6.27
P90 Ln(Weekly Wage)	7.01	6.97	7.12	6.92
Educational Background				
Share < High School	.15	.32	.08	.39
Share High School	.3	.39	.23	.4
Share Some College	.24	.2	.25	.17
Share College	.32	.09	.44	.04
Ethnicity				
Share Black	.1	.14	.08	.16
Share Hispanic	.06	.09	.04	.11
Industry Groups - Employment Shares				
Agriculture	0	.02	0	.02
Mining	0	.01	0	.01
Construction	.01	.08	.01	.12
Manufacturing	.12	.41	.12	.37
Transportation	.08	.14	.05	.13
Wholesale Trade	.06	.06	.06	.05
Retail Trade	.24	.07	.18	.1
Finance	.1	.02	.13	.02
Repair	.03	.04	.02	.03
Personal Services	.03	.01	.01	.03
Entertainment	.01	.01	.01	.01
Professional	.26	.09	.33	.08
Public Administration	.06	.04	.07	.03
Obs.	953,542	803,444	857,493	894,857
Weighted Obs.	16,682,180	15,647,427	16,564,660	16,331,087

NOTE.— Occupation-level summary statistics based on Software and Robot Exposure Quartiles. Exposure quartiles defined to include 1/4 of all workers in 1980 weighted sample. Sample restricted to employed individuals, age 16-64, in metropolitan areas that exist in both the 1980 and 2010 sample. Observations are weighted using ACS person weights adjusted to full-time full-year equivalence. Wages are winsorized at the bottom 1%. 1980 wages are adjusted for inflation using 2010 prices.Data: US Census & American Community Survey (1980-2010) (Ruggles et al., 2021). Exposure Measurement provided by Webb (2019).

The upper half of Figure 1 shows that software technology exposure is primarily high in male-dominated and mid to high-wage occupations. The software exposure to average wage relationship is similar for both men and women. [8] Regressing the logarithmic average wages on software exposure shows that the relationship between average wages and software exposure percentile is initially significant and positive but becomes insignificant, negative and small when controlling for demographics such as gender, age, ethnicity and education. The lower half of Figure 1 shows that the occupations exposed to robots are also male-dominated, but predominantly low-wage. The negative relationship between robot technology exposure rates and average wages is similar for both men and women. [9] Regressing the average wage on robot exposure percentiles shows that the relationship is significantly negative and sizeable, even when controlling for individual-level characteristics. [10]

Due to the definition of technology exposure, several occupations will be highly exposed to both robot and software technology simultaneously and could have exacerbating effects when combined. Furthermore, it could also be the case that changes in labor demand attributed to, for example, software technology exposure stem from these occupations' concurrent robot technology exposure and vice versa. The graph on the top left of Figure 2 shows the density of each combination of exposure quartiles. The density heat map shows that the majority of all observations occur along or close to the 45-degree line, meaning that exposure to software and robot technology is positively correlated. The top right graph shows the female share of workers in each combination of technology exposure quartile. The female share of workers is relatively high among occupations with below-median exposure in both technologies and occupations that are highly exposed to only of the technology categories considered here. Occupations with low robot exposure and high software exposure and vice versa are small and the overall relationship between exposure rates and female share is negative for both robot and software technology. The bottom left graph shows the average change in employment shares from 1980 to 2010 for each exposure quartile combination, weighted by the labor supply size of each quartile combination in 1980. Decreasing employment shares occur primarily in occupations with both high robot and high software exposure rates or high robot exposure overall. Increasing employment shares occur primarily in occupations with low levels of both types

[8] See Figure A2

[9] See Figure A2.

[10] The following descriptive regression was estimated for software and robot exposure separately:

$$ln(\text{Weekly wage})_i^{1980} = \alpha + \beta \text{Exp}_j^{\rightarrow 2018} + \mathbf{X}_i + \lambda_{k(i)} + \psi_{m(i)} + \varepsilon_i$$

. The specification includes controls for individual-level characteristics $\mathbf{X}_i$, such as age, educational background, ethnicity, and controls for industry, $\lambda_{k(i)}$, and metropolitan area of residence, $\psi_{m(i)}$. The complete results are presented in Tables A2 and A3 in the appendix.

Figure 1. Technology Exposure, Female Share and Average Wage - Occupation-Level

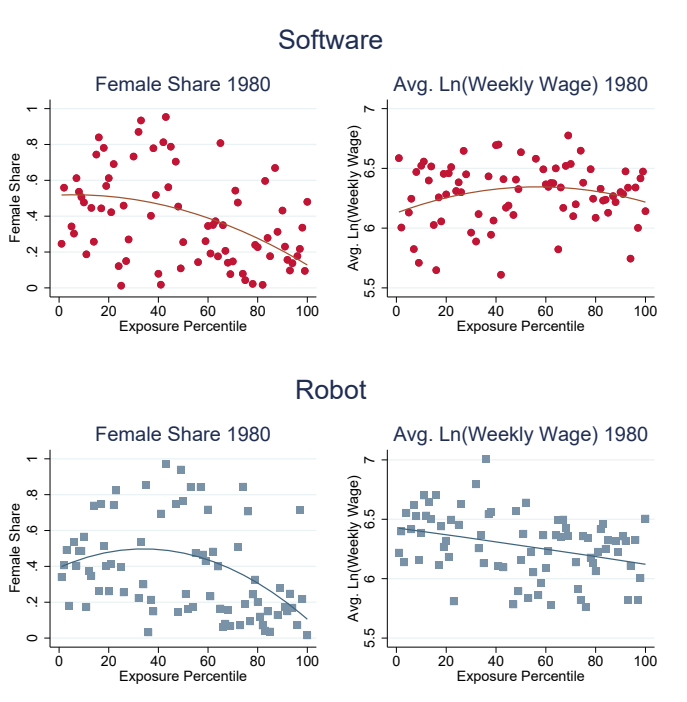

NOTE.— Binscatter of software and robot exposure percentile, reflecting future potential software or robot technology exposure and female share of workers and average wage in 1980 respectively. Female share and average wage is defined within each equally sized bin of exposure percentile, set at 100 bins.

Data: Weighted data from American Community Survey (1980), restricted to employed individuals ages 16-64 living in a metropolitan area (AMA) (Ruggles et al., 2021). Wages are winsorized at the bottom 1% and adjusted to 2010 price levels. Exposure measurement from Webb (2019).

of technology exposure. The bottom right graph shows the average log change in average weekly wages from 1980 to 2010 over each quartile combination, weighted by the 1980 labor supply size. Relatively smaller increases in the average wage are primarily prevalent among occupations with high robot exposure, regardless of software exposure rates. The overlap of exposure and the correlation with female employment shares must be taken into account when analyzing both changes in labor demand and the gender wage gap to accurately separate the effect of the two types of technology.

Figure 2. Heatmaps of Bivariate Technology Exposure - Occupation-Level

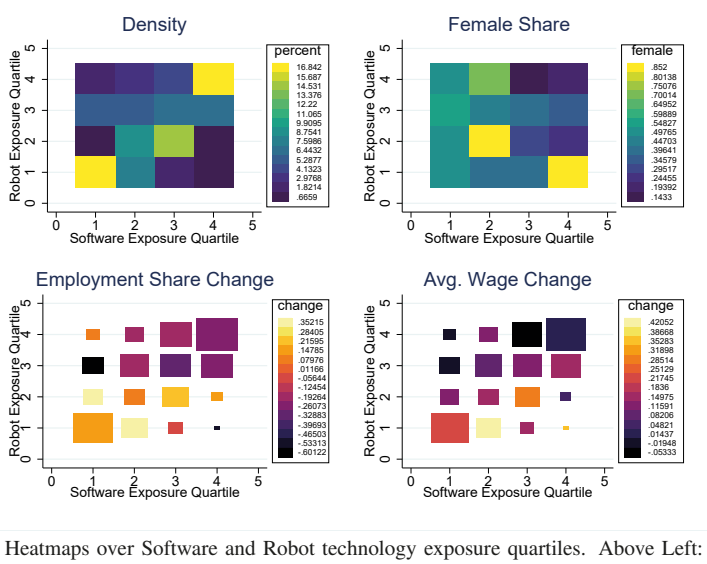

NOTE.— Heatmaps over Software and Robot technology exposure quartiles. Above Left: Density over Software and Robot exposure quartiles (1980). Above Right: Female share of workers in each technology exposure quartile combination (1980). Below Left: Arc Percentage Change in average employment shares from 1980-2010. Below Right: Log change in average weekly wages from 1980 to 2010. Exposure quartiles are defined to include 1/4 of all weighted observations in the 1980 sample. Data: Weighted data from American Community Survey (1980-2010), restricted to employed individuals ages 16-64 living in a metropolitan area (AMA) (Ruggles et al., 2021). Wages are winsorized at the bottom 1%. Exposure measurement from Webb (2019).

4.2 Technology Exposure and Labor Demand

I now explore how employment shares and average wages have developed over time depending on the Software or Robot exposure scores.

The upper left panel of Figure 3 shows how employment shares of occupations develop over time from 1980 to 2010 for each quartile of software exposure. The employment shares of each quartile are mechanically set around 25% [11] in 1980 according to the definition of the quartiles, but are allowed to develop differently over time. From 1980 to 2010, total employment shares of the most and mid-high exposed occupations decrease over time, especially among the most exposed occupations, whilst the employment shares of the least and mid-low exposed occupations increase. The upper right panel of Figure 3 shows how average weekly wages, all adjusted to 2010 price levels, develop over software exposure quartiles from 1980 to 2010. Average wages have increased in all exposure quartiles, but more so in the mid-low exposed occupations.

The lower left panel of Figure 3 shows the change in employment shares over robot exposure quartiles, defined in the same way as software exposure quartiles. The employment share of all quartiles has decreased from 1980 to 2010, except the least exposed quartile that has seen an increase in the relative employment share. The lower right panel of Figure 3 shows the changes in average wages over robot exposure quartiles and demonstrates that average wages are higher among the less exposed occupations in 1980 and this difference increases over time. Combined, it appears that the labor demand of highly exposed occupations, for both software and robot technology, has declined over time, whilst the labor demand has increased for less exposed occupations.

[11] Deviations from 25% come from all occupations with the same Software Exposure Percentile being delegated the same quartile assignment, making them somewhat uneven.

Figure 3. Change in Labor Demand over Exposure Quartiles - Occupation-Level

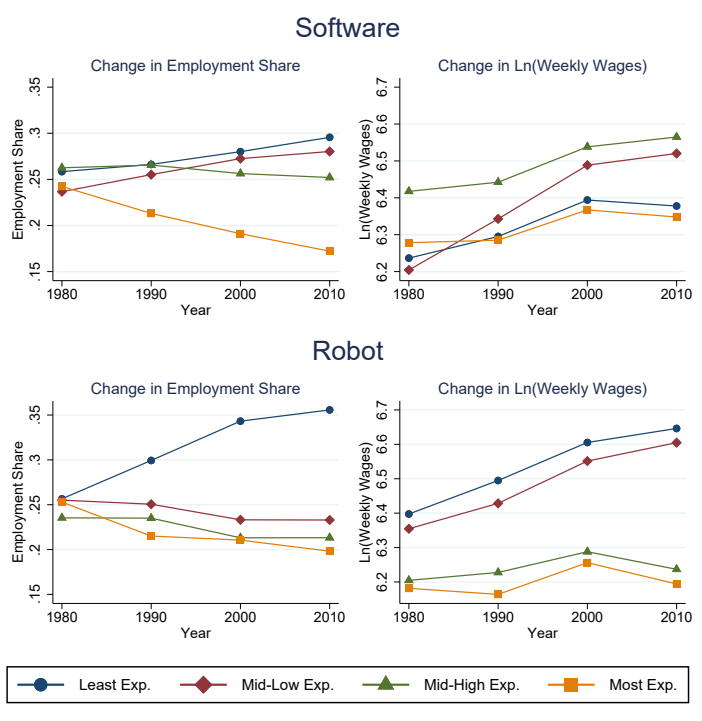

NOTE.— Above left: The change in relative employment shares for each software exposure quartile. Above right: The change in average Ln(Weekly Wages) for each software exposure quartile. Below left: The change in relative employment shares for each robot exposure quartile. Below right: The change in average Ln(Weekly Wages) for each robot exposure quartile. Data: Weighted data from American Community Survey (1980-2010), restricted to employed individuals ages 16-64 living in a metropolitan area (AMA) (Ruggles et al., 2021). Wages are winsorized at the bottom 1%. All wages are adjusted to 2010 levels. Quartiles are defined by occupation-level Software or Robot Exposure such that each quartile covers 1/4 of all workers in the 1980 sample, going from least to most exposed. Quartile assignments carry over to each year, regardless of employment share. Exposure measurement from Webb (2019).

Both men and women in declining technology-exposed occupations are affected by the decline in labor demand. The share of male and female labor supply employed in the most exposed occupations both decline significantly. The female share of workers does not appear to change significantly in software-exposed occupations, meaning that male and female workers are forced out at proportional rates. However, the female share of workers in robot-exposed occupations decreases disproportionately more over time. [12]

To further investigate the relationship between software and robot exposure and changes in labor demand, I estimate the following regression.

$$\Delta Y_j = \alpha + \beta \text{Exp}_j^{\rightarrow 2018} + \varepsilon_j \tag{3}$$

Where ΔY_j is either the long difference in total US employment share or average wage for a specific occupation j from 1980 to 2010. $\text{Exp}_j^{\rightarrow 2018}$ is the software or robot exposure percentile score, measuring the relative patenting activity up to and including 2018 for each occupation. To capture the entry and exit of certain occupation cells between 1980 and 2010, I opt for arc percentage changes, a symmetric measurement of growth where 2 denotes an entry and -2 denotes an exit (Davis et al., 1998). [13] Changes in average wages are defined as the log change in average wages, where 1980 wages are adjusted to 2010 price levels. The results for both software and robot technology are presented in Table 3.

Column 1 of the top panel of Table 3 shows that a higher software exposure percentile score is associated with a decrease in relative employment shares. This means that occupations with higher software technology exposure, between 1980 and 2010, declined in employment shares, indicating a decline in labor demand for these occupations relative to other, less exposed occupations. Column 2 shows a similar relationship between robot technology exposure and change in employment shares. Moving from the 25th to the 75th percentile of software or robot exposure is associated with a relative decline in employment growth of approximately 25%. [14] Column 3 controls for software and robot exposure, due to the overlap between exposure to robot and software technology. The coefficients on robot and software exposure are both still negative but much smaller and insignificant for software exposure. This is unsurprising considering the high correlation between software and robot

[12] For male and female share of labor supply, see Figure A4 and the top panel of A4. For the female share in occupation, see Figure A3 and the bottom panel of Table A4.

[13] Arc percentage change is defined as $\Delta X = \frac{X_{t_1} - X_{t_0}}{0.5(X_{t_0} + X_{t_1})}$

[14] Table A5 in the appendix shows a significant negative relationship between exposure and average wage growth as well as employment shares even when altering the end-year from 2010. The magnitude increases with time, demonstrating that the decline in labor demand does not come only from other labor market trends or shocks such as increased trade with China or the 2008 financial crisis.

exposure, and holding one fixed will remove much of the variation. [15] Column 4 controls for occupations' offshorability, to capture other labor market phenomena that could also explain decreases in employment shares depending on the task content of occupations. Offshorability is a measurement created by Autor and Dorn (2013) based on the same O*Net task taxonomy that is used as the basis for the measurement of technology exposure. Including this control does not affect the coefficients notably.

Column 1 of the bottom panel of Table 3 of shows a negative relationship between software exposure and the change in average wages. Moving from the 25th to the 75th percentile of software exposure is associated with a relative decline in wage growth of around 10%. Column 2 shows a significant negative relationship between robot exposure and the change in average wages. Moving from the 25th to the 75th percentile of robot exposure is associated with a relative decline in wage growth of around 19%. Column 3, similar to the top panel, takes into account the overlap in robot and software exposure. This reverses the sign on software exposure, but is now insignificant and greatly reduces the size of the coefficient on robot exposure. Column 4 controls for offshorability and does not have a visible effect on the coefficients. In general, high levels of software or robot exposure are associated with a lower average wage growth compared to occupations with lower levels of patenting activity.

[15] Table A6 shows that the negative relationship between exposure and employment shares is primarily derived from occupations with high robot *and* software exposure. The negative relationship between average wage growth and exposure is primarily derived among occupations with high robot exposure, regardless of software exposure.

Table 3. *Change in Labor Demand - Software and Robot Exposure (1980 - 2010)*

Δ Employment Share 1980-2010	(1)	(2)	(3)	(4)
Software Exposure until 2018	-0.518***		-0.206	-0.205
	(0.174)		(0.243)	(0.248)
Robot Exposure until 2018		-0.586***	-0.005**	-0.005**
		(0.160)	(0.002)	(0.002)
Offshorability				-0.040
				(0.046)
Constant	0.280***	0.298***	0.337***	0.344***
	(0.090)	(0.091)	(0.098)	(0.100)
Observations	324	324	324	324
R^2	0.05	0.07	0.08	0.09
Mean Dependent Variable	0.044	0.044	0.044	0.044
Δ Ln(Weekly Wages) 1980-2010	(1)	(2)	(3)	(4)
Software Exposure until 2018	-0.218***		0.068	0.068
	(0.066)		(0.070)	(0.070)
Robot Exposure until 2018		-0.379***	-0.004***	-0.004***
		(0.059)	(0.001)	(0.001)
Offshorability				-0.005
				(0.019)
Constant	0.268***	0.333***	0.320***	0.321***
	(0.035)	(0.035)	(0.035)	(0.036)
Observations	323	323	323	323
R^2	0.08	0.28	0.29	0.29
Mean Dependent Variable	0.169	0.169	0.169	0.169

NOTE.— Top Panel: Dependent variable is the long difference in employment shares from 1980 to 2010, calculated as arc percentage changes. Bottom Panel: Dependent variable is the long difference in Ln(Weekly Wage) from 1980 to 2010. Each observation is an occupation, defined by Dorn et al. (2009) occ1990dd. Software and Robot Exposure is measured at the occupation level and is defined as Software or Robot Exposure Percentile/100 for legibility reasons. All columns use observations weighted by occupation cell's labor supply, averaged between 1980 and 2010. Offshorability measures occupation-level eligibility of offshoring, created by Autor and Dorn (2013). Data: Weighted data from US Census & American Community Survey (1980-2010), restricted to employed inidividuals ages 16-64 living in a metropolitan area (AMA) (Ruggles et al., 2021). Wages are winsorized at the bottom 1%. Exposure measurement from Webb (2019). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

5 Technology Exposure across Cities

In this section, I explore how differences in initial employment shares in the most robot or software technology exposed occupations between cities may have affected the trajectory of the gender wage gap within these cities. The exposure measurement is only measured at the occupation level and does not vary over time nor over region. However, cities differ in their overall technology exposure since they differ in the employment share distribution of these occupations. It is reasonable to assume that, since software and robot technology exposure is primarily directed at male-dominated occupations and these occupations have subsequently seen a decline in labor demand, cities with overall more technology exposure (high exposure occupations have initially larger labor market shares) should have a sharper decline in the gender wage gap.

In an attempt to identify the relationship between technology exposure and changes in the gender wage gap, I aggregate my data from individual-level to city-level to now study how differences in city-level, defined as a metropolitan area in the American Community Survey, employment share distributions and their relationship to future potential technology exposure relates to the development of the gender wage gap over time. The trends shown in Figure 3, show that changes in employment shares occur in all quartiles of exposure, with occupations in the high and mid-high exposure quartiles declining in employment shares. Less exposed occupations have increased in employment shares. Based on this, I define city-level average software or robot exposure as the employment shares of the above median (3rd and 4th quartile) exposed occupations in 1980. This initial distribution of occupations, more or less exposed to robot or software technology in the future, sets the stage for how male and female workers will be affected relative each other and, in extention, how the gender wage gap may change over time. Figure A5 shows the distribution of city-level exposure for both software and robot technology and shows a relatively even spread in terms of having a majority of workers in either the most exposed or the less exposed occupations.

Using the employment share of high exposure (above median) occupations as the definition of city-level exposure, I divide the cities into two equally sized groups based on software exposure, to see in what characteristics these cities might differ, beyond differences in employment share distributions in more or less software technology exposed occupations. This is summarized in Table 4.

Table 4. *Summary Statistics - City-Level Means over Software Exposure (1980 - 2010)*

	1980		2010	
	Below Median	Above Median	Below Median	Above Median
Employment	396,579	250,005	689,436	372,837
Female Share	0.41	0.38	0.47	0.46
Ethnicity				
Share Black	0.10	0.07	0.13	0.09
Share Hispanic	0.06	0.04	0.14	0.11
Educational Background				
Share $<$ High School	0.19	0.23	0.07	0.07
Share High School	0.34	0.39	0.32	0.37
Share Some College	0.26	0.22	0.27	0.27
Share College	0.21	0.16	0.34	0.28
Weekly Wages & Gender Wage Gap				
Avg. Ln(Weekly Wage)	6.19	6.23	6.38	6.35
Avg. Female Ln(Weekly Wage)	6.11	6.10	6.50	6.43
Avg. Male Ln(Weekly Wage)	6.55	6.58	6.75	6.70
Avg. Gender Wage Gap	0.44	0.48	0.25	0.27
Software Exposure Quartile - Employment Shares				
Least Exposed Q1	0.27	0.23	0.29	0.27
Medium Low Exposed Q2	0.25	0.22	0.29	0.27
Medium High Exposed Q3	0.25	0.26	0.25	0.25
Most Exposed Q4	0.23	0.28	0.17	0.21
Robot Exposure Quartile - Employment Shares				
Least Exposed Q1	0.27	0.23	0.35	0.32
Medium Low Exposed Q2	0.25	0.22	0.23	0.21
Medium High Exposed Q3	0.24	0.25	0.22	0.24
Most Exposed Q4	0.24	0.30	0.20	0.24
Industry - Employment Shares				
Agriculture	0.01	0.02	0.02	0.02
Mining	0.01	0.01	0.01	0.01
Construction	0.06	0.06	0.05	0.05
Manufacturing	0.19	0.33	0.10	0.17
Transportation	0.09	0.08	0.07	0.07
Wholesale Trade	0.05	0.04	0.03	0.03
Retail Trade	0.15	0.14	0.17	0.17
Finance	0.07	0.05	0.07	0.06
Repair	0.03	0.03	0.06	0.05
Personal Services	0.03	0.02	0.03	0.02
Entertainment	0.01	0.01	0.02	0.01
Professional	0.22	0.18	0.31	0.29
Public Administration	0.08	0.05	0.07	0.05
Observations	239		239	

NOTE.— Summary statistics, city-level (Metropolitan Areas, AMA). Above and below median exposed cities defined by employment shares of the most exposed occupations. Sample restricted to employed individuals age 16-64, in metropolitan areas. Observations are weighted using ACS person weights adjusted to full-time full-year equivalence. Wages are winsorized at the bottom 1%. 1980 wages are adjusted for inflation using 2010 prices. Exposure quartiles are created separately for each technology category and are defined to include 1/4 of all workers in 1980 weighted sample, occupation-level quartile assignments carry over to 2010 sample. Data: US Census & American Community Survey (1980-2010) (Ruggles et al., 2021). Exposure Measurement provided by Webb (2019).

Cities with exposure rates above the median tend to be smaller on average than those below it. Among these cities, there is also a consistently lower proportion of workers with higher education qualifications over time. In 1980, cities with above-median exposure had higher average wages and a larger gender wage gap. However, by 2010, cities below the median exposure rate had surpassed them in average wages, likely due to slower wage growth in highly exposed occupations. Although the gender wage gap remains larger in the above-median cities, it has narrowed across all cities. Furthermore, these cities have a greater employment share in the manufacturing sector.

Figure 4 shows the relationship between city-level technology exposure in 1980 and the city's gender wage gap between 1950 and 2010. Point estimates are based on regressions of the gender wage gap on city-level technology exposure in 1980, separately for each sampled year between 1950 and 2010. The left panel of Figure 4 shows that before 1980 there was no significant relationship between city-level software exposure and the gender wage gap. In 1980, cities with larger employment shares of the most exposed occupations have on average higher gender wage gaps, but from 1980 and onward this gap diminishes. The right panel of Figure 4 shows that in 1970, relatively more robot-exposed cities had on average lower gender wage gaps, but this relationship disappears in 1980. From 1990 and onward, cities more exposed to robot technology have on average lower gender wage gaps.

Figure 4. City-Level Technology Exposure and Gender Wage Gap

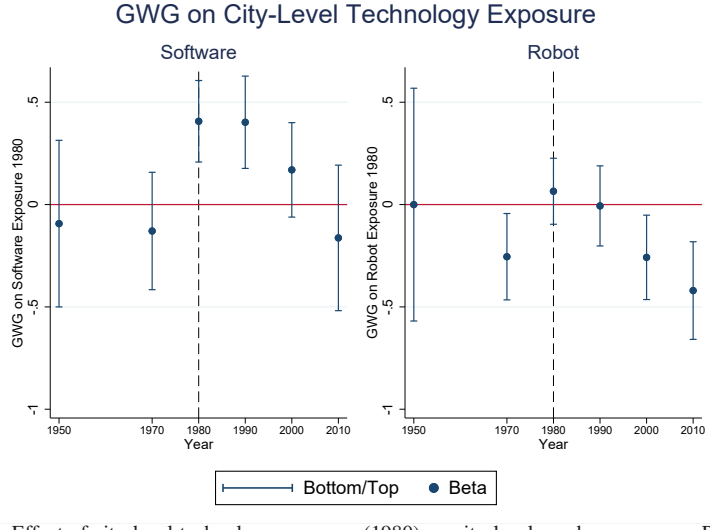

NOTE.— Effect of city-level technology exposure (1980) on city-level gender wage gap. Point estimates reflects the year-specific beta from the regression: $GWG_c = \alpha_0 + \beta^t Exp^{Technology}_{c,1980} + \varepsilon_c$ estimated separately for each year 1950-2010. Top/Bottom reflects the 95% confidence interval. Data: Weighted data from American Community Survey (1950-2010), restricted to employed individuals ages 16-64 living in a metropolitan area (AMA) (Ruggles et al., 2021). Each observation is a city that exists in the data in 1980 and 2010. Exposure measurement from Webb (2019).

To make the relationship between initial city-level technology exposure and the change in the gender wage gap more clear, I follow the cross-labor market empirical set-up by Acemoglu and Restrepo (2020) and estimate the following regression:

$$\Delta GWG_c = \alpha + \gamma \text{Exp}_{c,1980}^{Technology} + \beta X_c + \varepsilon_c \quad (4)$$

where ΔGWG_c is the nominal change in the gender wage gap for city c, between 1980 and 2010. $\text{Exp}_{c,1980}^{Technology}$ is the city-level average software or robot exposure rate in 1980, defined above, i.e., the employment shares of the above median exposed occupations. X_c includes city-level controls in the initial period, 1980. The identifying assumption is that cities are not differentially hit by other labor market shocks that should affect the trajectory of the gender wage gap, based on their initial employment shares in more or less exposed occupations. Furthermore, cities with initially larger employment shares of technology exposed occupations should not exhibit different trends in the gender wage gap, prior to large-scale technology adoption. The results of the regression equation 4 are presented in Table 5.

The top panel of Table 5 shows the estimates for various versions of equation 4, focusing on city-level software exposure. The overall impression is a robust, negative, relationship between city-level software exposure in 1980 and the change in the gender wage gap. This means that the gender wage gap declined more in cities with initially larger employment shares of software-exposed occupations.

The first column shows the raw association, and the following columns gradually add more controls. Column 2 adds controls for city-level demographics, measured in 1980. Column 3 substitutes the dependent variable, the nominal change in the gender wage gap, with the nominal change in the residual gender wage gap. Column 4 controls for pre-existing trends in the form of the change in the gender wage gap from 1950 to 1970. The number of observations declines because fewer cities had metropolitan status in 1950 to 1970 compared to 1980 to 2010. Column 5 shows the baseline model with residual gender wage gap for all cities with metropolitan status between 1950 and 2010. [16] Column 6 controls for simultaneous city-level robot exposure. Column 7 removes the pretrends in the gender wage gap to increase the number of observations.

The overall results are robust and show that moving from the 25th to the 75th percentile of city-level software exposure in 1980 is associated with an increased decline in the gender wage gap of around 3 to 3.5 percentage points. This accounts for 15-18% of the total average decline in the sampled cities. What seems to matter somewhat for the magnitudes, beyond city-level demographics, is primarily the control for preexisting trends in the gender wage gap. The results show that the decline in the gender wage gap of software-exposed

[16] Table A12 recreates Table 5, for the 106 cities sampled in all years.

Table 5. *Main Results: Gender Wage Gap and City-Level Technology Exposure (1980 - 2010)*

	Δ GWG		Δ Residual GWG				
Software Exposure	(1)	(2)	(3)	(4)	(5)	(6)	(7)
$Exp_{c,1980}^{Software}$	-0.601***	-0.498***	-0.489***	-0.543***	-0.537***	-0.519***	-0.546***
	(0.160)	(0.133)	(0.088)	(0.117)	(0.116)	(0.150)	(0.105)
Δ Res. GWG 1950-1970				0.010		0.011	
				(0.041)		(0.041)	
$Exp_{c,1980}^{Robot}$						-0.055	0.123
						(0.243)	(0.145)
Constant	0.108	-0.122	-0.176***	-0.167**	-0.172**	-0.143	-0.223**
	(0.082)	(0.095)	(0.066)	(0.084)	(0.084)	(0.139)	(0.087)
Observations	239	239	239	106	106	106	239
R^2	0.13	0.31	0.45	0.58	0.58	0.58	0.45
Mean Dependent Variable	-0.192	-0.192	-0.171	-0.171	-0.171	-0.171	-0.171
Demographic Controls		✓	✓	✓	✓	✓	✓
Weighted Obs.	✓	✓	✓	✓	✓	✓	✓

	Δ GWG		Δ Residual GWG				
Robot Exposure	(1)	(2)	(3)	(4)	(5)	(6)	(7)
$Exp_{c,1980}^{Robot}$	-0.410***	-0.464***	-0.298**	-0.526***	-0.531***	-0.055	0.123
	(0.122)	(0.162)	(0.126)	(0.188)	(0.184)	(0.243)	(0.145)
Δ Res. GWG 1950-1970				-0.008		0.011	
				(0.043)		(0.041)	
$Exp_{c,1980}^{Software}$						-0.519***	-0.546***
						(0.150)	(0.105)
Constant	0.007	-0.100	-0.253***	-0.106	-0.102	-0.143	-0.223**
	(0.062)	(0.119)	(0.087)	(0.132)	(0.129)	(0.139)	(0.087)
Observations	239	239	239	106	106	106	239
R^2	0.10	0.29	0.40	0.53	0.53	0.58	0.45
Mean Dependent Variable	-0.192	-0.192	-0.171	-0.171	-0.171	-0.171	-0.171
Demographic Controls		✓	✓	✓	✓	✓	✓
Weighted Obs.	✓	✓	✓	✓	✓	✓	✓

NOTE.— Dependent variable is the nominal difference in a city's gender wage gap from 1980 to 2010. Residualized gender wage gap uses wages residualized on age, ethnicity and educational background. Each observation is a city in the final sample. Software ($Exp_{c,1980}^{Software}$) and Robot ($Exp_{c,1980}^{Robot}$) Exposure is defined as the employment share of the above median exposed occupations within a city 1980. Observations are weighted by 1980 labor supply size. City-level demographics include ethnicity and educational background. Number of observations decrease in column 3 due to the smaller number of cities with metropolitan status in 1950-1970. Data: Weighted data from American Community Survey (1950-2010), restricted to employed inidividuals ages 16-64 living in a metropolitan area (AMA) (Ruggles et al., 2021). Wages are winsorized at the bottom 1%. 1980 wages are inflated to 2010 levels. Exposure measurement from Webb (2019). * p $<$ 0.10, ** p $<$ 0.05, *** $p < 0.01$

cities is not driven by an already steeper decline prior to software technology adoption. [17] Switching from the regular gender wage gap to the residual gender wage gap did not change the magnitudes of the estimates noticeably. The negative association between software exposure and the gender wage gap does not appear to be driven by the selection of workers in and out of employment. [18] As a robustness check, I redefine city-level software exposure as the share of workers in the most exposed quartile, rather than the above median exposed, and find that the results and conclusions are robust to alternative definitions in Table A13. Redefining city-level exposure as the continuous measurement employment-weighted percentile score also does not affect the conclusions, as seen in Table A14.

I perform the same exercise as the top panel of Table 5 in the bottom panel, substituting Software exposure for Robot technology exposure. City-level exposure is defined in the same way as before, i.e., the employment shares of above median exposed occupations in 1980, now using robot patenting activity as the exposure measurement. The results show that the gender wage gap declined more in cities with initially higher employment shares of robot-exposed occupations. However, this relationship is not robust for controls for simultaneous software exposure. Furthermore, substituting the change in the gender wage gap with the change in the residual gender wage gap matters for the magnitude of the estimates. As discussed in section 2, the increased supply of workers in less exposed, low-wage occupations should drive down wages, for both new and incumbent workers. However, Table A7 shows that high rates of robot exposure reduces the male employment-to-population ratio, positing that male workers, displaced by robot technology, entered unemployment to a large extent. This changes the demographic make-up of the male labor force, in terms of educational background, for example, and this matters for the gender wage gap. In addition to this, the city-level demographic controls includes educational background in the initial period. [19] Overall, high levels of both software and robot exposure at the city level are associated with a more rapid decline in the gender wage gap between 1980 and 2010. [20]

[17] The top panel of Table A11 shows that software-exposed cities had, if anything, increasing gender wage gaps between 1950 and 1970 compared to less exposed cities.

[18] The top panel of Table A7 shows that software-exposed cities declined in the employment-to-population ratio from 1980 to 2010. The middle and bottom panels show that male and female workers appear differently affected by this and to capture demographic changes, the residual gender wage gap is preferable.

[19] Table A8 shows no significant relationship between city-level software exposure and the share of college educated cities, but robot-exposed cities demonstrated a larger decline in the share of college educated cities, both overall and for female workers specifically.

[20] Table A15 shows that for software exposure, between 1980 to 1990 and 2000 respectively, the gender wage gap did decline more in exposed cities, however, not significantly. This is in line with the results in table A5, showing that although employment shares and average wages decline early (1980-1990), magnitudes increase over time.

6 The future of the Gender Wage Gap and AI Technology

I am able to show that the gender wage gap, between 1980 and 2010, declined more in cities with higher employment shares of robot- and software-exposed occupations. Although early in its development and implementation, artificial intelligence has already shown potential in both augmenting and replacing labor. During the early 2000s, the capabilities of AI were largely concentrated on single tasks.[21] From the late 2000s and onward, the onset of big data, deep learning, and Artificial General Intelligence (AGI) quickly expanded the applications of AI in the economy (Manyika et al., 2011). In this section, I perform an exploratory analysis of the role of AI in the evolution of the gender wage gap moving forward. I will characterize the group of workers most exposed to AI patenting activity in 2019 and discuss how it will determine the gender wage gap in the coming decades, assuming the same trends and mechanisms as with robot and software technology.[22]

The upper left panel of Figure 5 shows that, similar to both software and robot exposure in 1980, AI exposure is highest in primarily male-dominated occupations. The top right panel of Figure 5 shows that AI exposure is high in predominantly high-wage occupations. Table 6 shows sample means for each quartile of AI exposure. AI-exposed occupations are, beyond male-dominated and high-wage, also more common among college-educated workers. Furthermore, AI-exposed occupations occur more in the professional industry category, more so than robot- and software-exposed occupations. Regressing the average wage of occupations on the AI exposure rate shows that the relationship is sizeable and positive, even when controlling for workers characteristics such as gender, age, educational background, and ethnicity. It also shows that the positive relationship is somewhat stronger for female workers. [23] Assuming the same replacement effects and decreases in labor demand, as with robot and software technology, AI-induced labor displacement will affect mainly male workers. Compared to software-exposed occupations, and especially robot-exposed occupations, AI-exposed occupations have much higher average wages, and this should exacerbate the effect on the gender wage gap. Table 7 shows the change in the gender wage gap between 2000 and 2019 on

[21] Deep Blue beating world chess champion Garry Kasparov (McCorduck and Cfe, 2004), IBMs Watson beating Jeapordy! champions Brad Rutter and Ken Jennings (Ferrucci, 2012)

[22] Although more recent data is available from IPUMS ACS, I limit the study period to 2019 to avoid confounding labor market effects of the Covid-19 pandemic.

[23] The following descriptive regression was estimated for AI exposure:

$$ln(\text{Weekly wage})_i^{2000} = \alpha + \beta \text{Exp}_j^{\rightarrow 2018} + \mathbf{X}_i + \lambda_{k(i)} + \psi_{m(i)} + \varepsilon_i$$

The specification includes controls for individual-level characteristics, $\mathbf{X}_i$, such as age, educational background, ethnicity and controls for industry, $\lambda_{k(i)}$, and metropolitan area of residence, $\psi_{m(i)}$. Full results are presented in Table A16.

Figure 5. Labor Characteristics and Software Exposure over AI Exposure

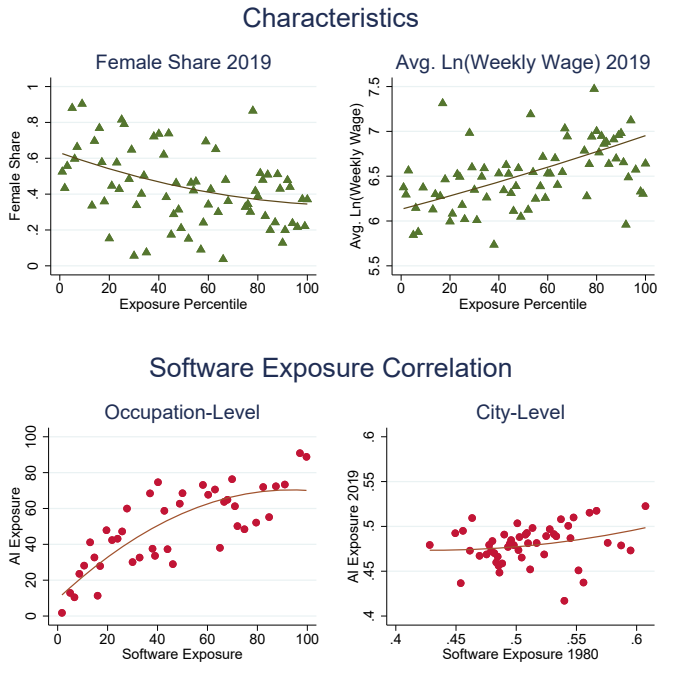

NOTE.— Top: Binscatter of AI exposure percentiles, reflecting future potential AI technology exposure (until 2018) and female share of workers and average wage in 2019 respectively. Female share and average wage is defined within each equally sized bin of exposure percentile, set at 100 bins. Bottom: Graphs showing exposure correlation between AI and Software technology at the occupation and city-level, respectively. City-level AI and Software exposure is defined as the employment share of above median exposed occupations, in 2019 and 1980 respectively. Data: Weighted data from American Community Survey (2019), restricted to employed individuals ages 16-64 living in a metropolitan area (AMA) (Ruggles et al., 2021). Wages are adjusted to full-time full-year equivalence, winsorized at the bottom 1% and adjusted to 2010 price levels. Exposure measurement from Webb (2019)

city-level AI exposure, measured in 2000. The estimates show that the gender wage gap declined more in AI-exposed cities, but the coefficients are not statistically significant. However, it is likely too early to see the full effects of AI exposure on the gender wage gap.

The bottom left panel of Figure 5 shows a positive relationship between occupation-level software and AI exposure. This means that several occupations that were most exposed during the onset of computerization are now affected by the emergence of AI technology. The bottom right panel shows a weak, yet positive, relationship between city-level software exposure in 1980 and city-level AI exposure in 2019. This means that cities that already showed larger declines in the gender wage gap, due to changes in labor demand in-

duced by software technology, will show further declines compared to less exposed occupations.

Table 6. *Summary Statistics - Occupation-Level Means over AI Exposure Quartiles*

	Quartile 1	Quartile 2	Quartile 3	Quartile 4
Female Share	.6	.44	.39	.39
Avg. Age	39.63	40.87	41.87	41.25
Weekly Wage				
Avg. Ln(Weekly Wage)	6.22	6.4	6.69	6.81
P10 Ln(Weekly Wage)	5.41	5.58	5.81	5.92
Median Ln(Weekly Wage)	6.2	6.39	6.65	6.83
P90 Ln(Weekly Wage)	7.06	7.22	7.59	7.61
Educational Background				
Share < High School	.07	.08	.05	.04
Share High School	.36	.35	.27	.2
Share Some College	.29	.22	.22	.21
Share College	.27	.36	.45	.55
Ethnicity				
Share Black	.16	.15	.11	.1
Share Hispanic	.24	.24	.19	.16
Industry Groups - Employment Shares				
Agriculture	.01	.01	.03	.02
Mining	0	0	.01	.01
Construction	.03	.09	.09	.06
Manufacturing	.05	.08	.16	.15
Transportation	.06	.14	.06	.07
Wholesale Trade	.04	.03	.03	.02
Retail Trade	.29	.14	.14	.05
Finance	.03	.08	.09	.09
Repair	.05	.07	.07	.12
Personal Services	.04	.02	.02	.01
Entertainment	.02	.02	.02	.01
Professional	.36	.28	.25	.3
Public Administration	.03	.04	.05	.08
Obs.	280,630	245,997	220,368	241,004
Weighted Obs.	30,414,436	29,642,088	28,464,716	29,803,016

NOTE.— Occupation-level summary statistics based on AI Exposure Quartiles. Exposure quartiles defined to include 1/4 of all workers in 2019 weighted sample. Sample restricted to employed individuals, age 16-64, in metropolitan areas that exist in both the 1980 and 2010 sample. Observations are weighted using ACS person weights adjusted to full-time full-year equivalence. Wages are winsorized at the bottom 1%. Data: US Census & American Community Survey (2019) (Ruggles et al., 2021). Exposure Measurement provided by Webb (2019).

Table 7. *Gender Wage Gap and city-level AI Exposure (2000-2019)*

	Δ GWG		Δ Residual GWG				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
$Exp^{AI}_{c,2000}$	-0.274	-0.275	-0.290*	-0.263	-0.262	-0.027	-0.063
	(0.219)	(0.206)	(0.158)	(0.201)	(0.187)	(0.234)	(0.173)
Δ Residual GWG 1970-1990				0.001		0.002	
				(0.079)		(0.080)	
$Exp^{Software}_{c,1980}$						-0.171	-0.186
						(0.207)	(0.176)
$Exp^{Robot}_{c,1980}$						-0.277	-0.223
						(0.197)	(0.142)
Constant	0.088	-0.006	0.056	-0.039	0.091	0.055	0.243**
	(0.108)	(0.122)	(0.079)	(0.154)	(0.222)	(0.175)	(0.120)
Observations	174	174	174	101	101	101	174
R^2	0.02	0.17	0.27	0.32	0.32	0.36	0.31
Mean Dependent Variable	-0.044	-0.044	-0.009	-0.008	-0.008	-0.008	-0.009
Demographic Controls		✓	✓	✓	✓	✓	✓
Weighted Obs.	✓	✓	✓	✓	✓	✓	✓

NOTE.— Dependent variable is the nominal difference in a city's gender wage gap from 2000 to 2019. Each observation is a city in the final sample. AI Exposure ($Exp^{AI}_{c,2000}$) is defined as the employment share of the above median exposed occupations within a city 2000. Software and Robot Exposure is defined as the employment share of above median software or robot exposed occupations. Column 6 use prior city-level software and robot exposure, measured in 1980. Observations are weighted by 2000 labor supply size. City-level demographics include ethnicity and educational background. Number of observations decrease in columns 4-6 due to the smaller number of cities with metropolitan status in 1970-1990. Column 5 holds the number of cities fixed from 1970 to 2019. Data: Weighted data from American Community Survey (1970-2019), restricted to employed inidividuals ages 16-64 living in a metropolitan area (AMA) (Ruggles et al., 2021). Wages are winsorized at the bottom 1%. 1980 wages are inflated to 2010 levels. Exposure measurement from Webb (2019).
* p < 0.10, ** p < 0.05, *** p < 0.01

7 Conclusions

The adoption of robot and software technology in the past decades has played a significant role in reshaping the landscape of the labor market. Although much work has been done to characterize the types of workers, industries and labor markets most affected by it, little has been done to identify the role technological innovation has played in the development of the gender wage gap. In this paper, I have shown that the gender wage gap declined more in US cities that were more exposed to software and robot technology, through larger employment shares in occupations most exposed to robot or software patenting activity. I posit that this is due to technology's ability to replace human labor in performing tasks and that this has had a negative impact on the labor demand in primarily male-dominated occupations. However, the relationship between city-level robot exposure and the decline in the gender wage gap does not hold when controlling for simultaneous software technology exposure and I attribute this partly to the overlap in robot and software technology exposure,

both at the occupation- and the city-level, and partly to robot exposure being primarily high in low-wage occupations.

Understanding how initial occupational distribution and gender segregation at the local labor market level sets the stage for how technological advancements have and will effect the gender wage gap is key to understanding what the large-scale adoption of AI-related technology will mean for the gender wage gap in the future.

References

Acemoglu, Daron (2003). "Technology and inequality". *Technology and Inequality*.

Acemoglu, Daron and David Autor (2011). "Skills, tasks and technologies: Implications for employment and earnings". *Handbook of labor economics*. Vol. 4. Elsevier, 1043–1171.

Acemoglu, Daron and Pascual Restrepo (2018). "Artificial intelligence, automation, and work". *The economics of artificial intelligence: An agenda*. University of Chicago Press, 197–236.

Acemoglu, Daron and Pascual Restrepo (2020). "Robots and jobs: Evidence from US labor markets". *Journal of Political Economy* 128.6, 2188–2244.

Autor, D and David Dorn (2013). "The growth of low-skill service jobs and the polarization of the US labor market". *American economic review* 103.5, 1553–97.

Autor, David, Caroline Chin, Anna Salomons, and Bryan Seegmiller (2024). "New frontiers: The origins and content of new work, 1940–2018". *The Quarterly Journal of Economics*, qjae008.

Autor, David H, Lawrence F Katz, and Alan B Krueger (1998). "Computing inequality: Have computers changed the labor market?" *The Quarterly journal of economics* 113.4, 1169–1213.

Autor, David H, Frank Levy, and Richard J Murnane (2003). "The skill content of recent technological change: An empirical exploration". *The Quarterly journal of economics* 118.4, 1279–1333.

Beaudry, Paul, Mark Doms, and Ethan Lewis (2010). "Should the personal computer be considered a technological revolution? Evidence from US metropolitan areas". *Journal of political Economy* 118.5, 988–1036.

Beaudry, Paul and Ethan Lewis (2014). "Do male-female wage differentials reflect differences in the return to skill? Cross-city evidence from 1980-2000". *American Economic Journal: Applied Economics* 6.2, 178–94.

Becker, Gary S, William HJ Hubbard, and Kevin M Murphy (2010). "Explaining the worldwide boom in higher education of women". *Journal of Human Capital* 4.3, 203–241.

Black, Sandra E and Alexandra Spitz-Oener (2010). "Explaining women's success: Technological change and the skill content of women's work". *The Review of Economics and Statistics* 92.1, 187–194.

Blackburn, Robert M, Jennifer Jarman, and Bradley Brooks (2000). "The puzzle of gender segregation and inequality: A cross-national analysis". *European Sociological Review* 16.2, 119–135.

Blau, Francine D, Peter Brummund, and Albert Yung-Hsu Liu (2013). "Trends in occupational segregation by gender 1970–2009: Adjusting for the impact

of changes in the occupational coding system". *Demography* 50.2, 471–492.

Blau, Francine D and Lawrence M Kahn (1997). "Swimming upstream: Trends in the gender wage differential in the 1980s". *Journal of labor Economics* 15.1, Part 1, 1–42.

Blau, Francine D and Lawrence M Kahn (2006). "The US gender pay gap in the 1990s: Slowing convergence". *Ilr Review* 60.1, 45–66.

Blau, Francine D and Lawrence M Kahn (2013). "Female labor supply: Why is the United States falling behind?" *American Economic Review* 103.3, 251–256.

Blau, Francine D and Lawrence M Kahn (2017). "The gender wage gap: Extent, trends, and explanations". *Journal of economic literature* 55.3, 789–865.

Cortes, Guido Matias (2016). "Where have the middle-wage workers gone? A study of polarization using panel data". *Journal of Labor Economics* 34.1, 63–105.

Cortes, Guido Matias, Nir Jaimovich, and Henry E Siu (2018). "The "end of men" and rise of women in the high-skilled labor market". Tech. rep. National Bureau of Economic Research.

Cortes, Guido Matias, Ana Oliveira, and Anna Salomons (2020). "Do technological advances reduce the gender wage gap?" *Oxford Review of Economic Policy* 36.4, 903–924.

Davis, Steven J, John C Haltiwanger, Scott Schuh, et al. (1998). "Job creation and destruction". *MIT Press Books* 1.

Dechezleprêtre, Antoine, David Hémous, Morten Olsen, and Carlo Zanella (2021). "Induced automation: Evidence from firm-level patent data". *University of Zurich, Department of Economics, Working Paper* 384.

Dorn, David et al. (2009). "This job is "getting old": Measuring changes in job opportunities using occupational age structure". *American Economic Review* 99.2, 45–51.

Ferrucci, David A (2012). "Introduction to "This is Watson"". *IBM Journal of Research and Development* 56.3.4, 1–1.

Goldin, Claudia (2006). "The quiet revolution that transformed women's employment, education, and family". *American economic review* 96.2, 1–21.

Goldin, Claudia, Lawrence F Katz, and Ilyana Kuziemko (2006). "The homecoming of American college women: The reversal of the college gender gap". *Journal of Economic perspectives* 20.4, 133–156.

Goos, Maarten and Alan Manning (2007). "Lousy and lovely jobs: The rising polarization of work in Britain". *The review of economics and statistics* 89.1, 118–133.

Hegewisch, Ariane and Heidi Hartmann (2014). "Occupational segregation and the gender wage gap: A job half done". Tech. rep. Institute for Women's Policy Research.

Holzer, Harry and David Neumark (2000). "Assessing affirmative action". *Journal of Economic literature* 38.3, 483–568.

Kogan, Leonid, Dimitris Papanikolaou, Lawrence DW Schmidt, and Bryan Seegmiller (2021). *Technology-skill complementarity and labor displacement: Evidence from linking two centuries of patents with occupations*. National Bureau of Economic Research.

Lindley, Joanne (2012). "The gender dimension of technical change and the role of task inputs". *Labour Economics* 19.4, 516–526.

Mann, Katja and Lukas Püttmann (2023). "Benign effects of automation: New evidence from patent texts". *Review of Economics and Statistics* 105.3, 562–579.

Manyika, James, Michael Chui, Brad Brown, Jacques Bughin, Richard Dobbs, Charles Roxburgh, Angela Hung Byers, et al. (2011). *Big data: The next frontier for innovation, competition, and productivity*. McKinsey Global Institute.

McCorduck, Pamela and Cli Cfe (2004). *Machines who think: A personal inquiry into the history and prospects of artificial intelligence*. CRC Press.

Ruggles, Steven, Sarah Flood, Sophia Foster, Ronald Goeken, Megan Schouweiler Jose Pacas, and Matthew Sobek (2021). "PUMS USA: Version 11.0 [dataset]".

Spitz-Oener, Alexandra (2006). "Technical change, job tasks, and rising educational demands: Looking outside the wage structure". *Journal of labor economics* 24.2, 235–270.

Webb, Michael (2019). "The impact of artificial intelligence on the labor market". *Available at SSRN 3482150*.

Appendix A: Additional figures and tables

Figure A1. Histogram of Male and Female Weekly Wages 1980 - 2010

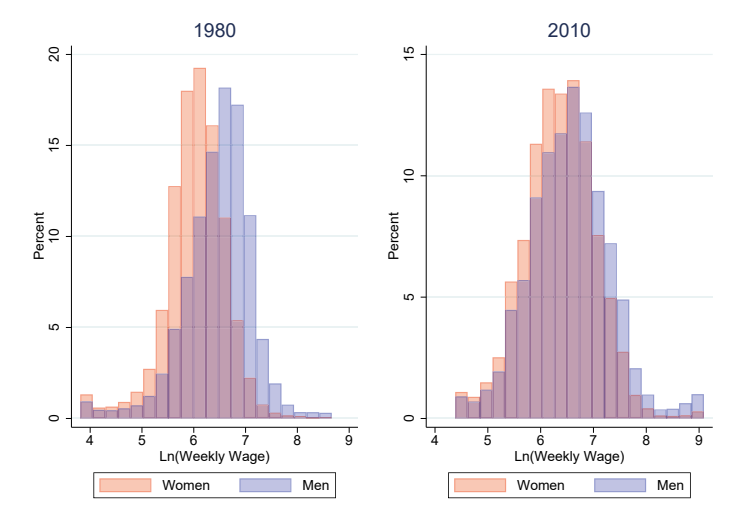

NOTE.— Histogram of male and female ln(Weekly wages) from 1980 and 2010 respectively. Data: Weighted data from American Community Survey (1980-2010), restricted to employed individuals ages 16-64 living in a metropolitan area (AMA) (Ruggles et al., 2021). Wages are winsorized at the bottom 1%. 1980 wages are adjusted for inflation using 2010 prices.

Table A1. *Most and Least Exposed Occupations for each technology category*

Technology	Most Exposed Occupations	Least Exposed Occupations
Software	Broadcast equipment operators	Podiatrists
	Pest control occupations	Subject instructors, college
	Elevator installers and repairers	Art/entertainment performers and related occupations
	Water and sewage treatment plant operators	Mail carriers for postal service
	Power plant operators	Barbers
Robot	Industrial truck and tractor operators	Art/entertainment performers and related occupations
	Operating engineers of construction equipment	Mail carriers for postal service
	Janitors	Correspondence and order clerks
	Elevator installers and repairers	Clergy
	Housekeepers, maids, butlers, and cleaners	Eligibility clerks for government prog., social welfare
AI	Atmospheric and space scientists	Art/entertainment performers and related occupations
	Metallurgical and materials engineers	Mail carriers for postal service
	Purchasing agents and buyers of farm products	Subject instructors, college
	Marine engineers and naval architects	Funeral directors
	Clinical laboratory technologies and technicians	Animal caretakers, except farm

NOTE.— Most and least exposed occupations for each type of technology, Software, Robot and AI. Exposure measurement from Webb (2019).

Table A2. *Ln(Weekly Wages) and Occupation-Level Software Exposure Percentile*

Ln(Weekly Wage)$_{i,1980}$	(1)	(2)	(3)	(4)	(5)
Software Exposure until 2018	0.107***	-0.065***	-0.095***	-0.038***	-0.037***
	(0.001)	(0.001)	(0.002)	(0.001)	(0.001)
Female		-0.435***	-0.474***	-0.321***	-0.322***
		(0.001)	(0.001)	(0.001)	(0.001)
Female x Software Exp.			0.084***		
			(0.003)		
Average Age				0.013***	0.013***
				(0.000)	(0.000)
Share Black				-0.109***	-0.121***
				(0.001)	(0.001)
Share Hispanic				-0.091***	-0.106***
				(0.001)	(0.002)
Share High School				0.188***	0.184***
				(0.001)	(0.001)
Share Some College				0.273***	0.265***
				(0.001)	(0.001)
Share College				0.558***	0.543***
				(0.001)	(0.001)
Observations	3,413,606	3,413,606	3,413,606	3,413,606	3,413,606
R^2	0.00	0.10	0.10	0.32	0.33
Weighted Obs.	64,559,931	64,559,931	64,559,931	64,559,931	64,559,931
Industry Dummies				✓	✓
AMA Dummies					✓

NOTE.— Dependent variable is Ln(Weekly Wage). Software Exposure is defined as Software Exposure Percentile/100 for legibility reasons. Industry dummies include all industries as defined by 'ind1990' in ACS. AMA dummies include all metropolitan areas that exist in the data both in 1980 and 2010. The reference variable for ethnicity is White. The reference variable for Educational Background is 'Less than High School'. Data: Weighted data from American Community Survey (ACS, 1980), restricted to employed inidividuals ages 16-64 living in a metropolitan area (AMA) (Ruggles et al., 2021). Wages are winsorized at the bottom 1%. Exposure measurement from Webb (2019).
* p $<$ 0.10, ** p $<$ 0.05, *** $p < 0.01$

Table A3. *Ln(Weekly Wages) and Occupation-Level Robot Exposure Percentile*

Ln(Weekly Wage)$_{i,1980}$	(1)	(2)	(3)	(4)	(5)
Robot Exposure until 2018	-0.306***	-0.450***	-0.452***	-0.230***	-0.226***
	(0.001)	(0.001)	(0.002)	(0.001)	(0.001)
Female		-0.476***	-0.479***	-0.335***	-0.336***
		(0.001)	(0.001)	(0.001)	(0.001)
Female x Robot Exp.			0.008***		
			(0.003)		
Average Age				0.013***	0.013***
				(0.000)	(0.000)
Share Black				-0.096***	-0.106***
				(0.001)	(0.001)
Share Hispanic				-0.083***	-0.096***
				(0.001)	(0.002)
Share High School				0.168***	0.165***
				(0.001)	(0.001)
Share Some College				0.236***	0.228***
				(0.001)	(0.001)
Share College				0.490***	0.477***
				(0.001)	(0.001)
Observations	3,413,606	3,413,606	3,413,606	3,413,606	3,413,606
R^2	0.02	0.13	0.13	0.32	0.34
Weighted Obs.	64,559,931	64,559,931	64,559,931	64,559,931	64,559,931
Industry Dummies				✓	✓
AMA Dummies					✓

NOTE.— Dependent variable is Ln(Weekly Wage). Robot Exposure is defined as Robot Exposure Percentile/100 for legibility reasons. Industry dummies include all industries as defined by 'ind1990' in ACS. AMA dummies include all metropolitan areas that exist in the data both in 1980 and 2010. The reference variable for ethnicity is White. The reference variable for Educational Background is 'Less than High School'. Data: Weighted data from American Community Survey (ACS, 1980), restricted to employed inidividuals ages 16-64 living in a metropolitan area (AMA) (Ruggles et al., 2021). Wages are winsorized at the bottom 1%. Exposure measurement from Webb (2019).
* p $<$ 0.10, ** p $<$ 0.05, *** $p < 0.01$

Figure A2. Technology Exposure, Male and Female Average Wage - Occupation-Level

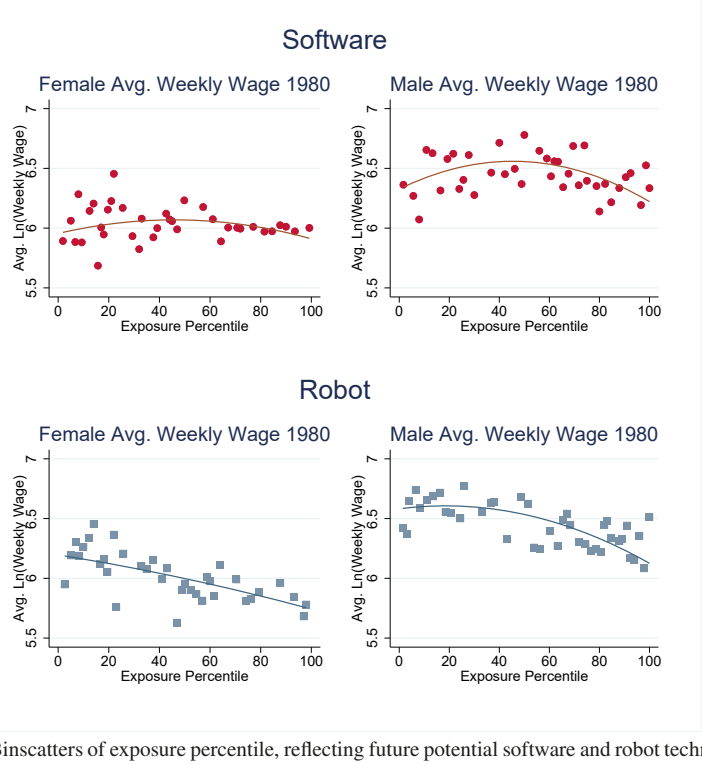

NOTE.— Binscatters of exposure percentile, reflecting future potential software and robot technology exposure, and female and male average wages in 1980 respectively. Male and female average wages are defined within each equally sized bin of exposure percentile, set at 50 bins. Data: Weighted data from American Community Survey (1980), restricted to employed individuals ages 16-64 living in a metropolitan area (AMA) (Ruggles et al., 2021). Wages are winsorized at the bottom 1% and adjusted to 2010 price levels. Exposure measurement from Webb (2019).

Figure A3. Change in Female Share over Technology Exposure - Occupation-Level

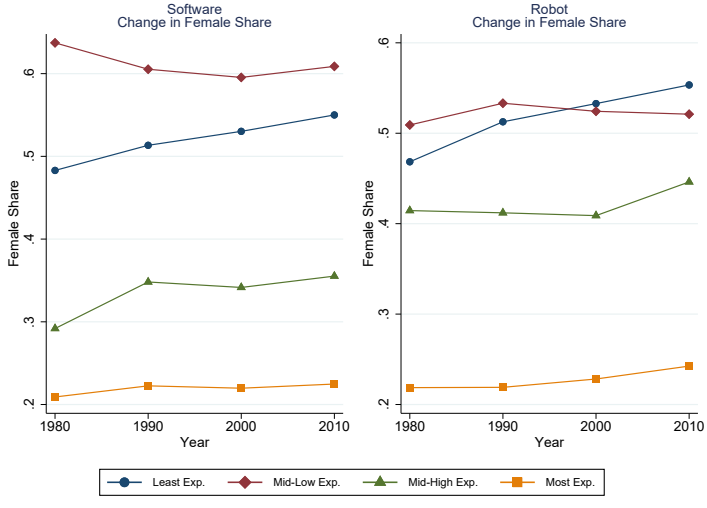

NOTE.— Right: Change in the female share of workers within each software exposure quartile from 1980 to 2010. Left: Change in the female share of workers within each robot exposure quartile from 1980 to 2010. Quartiles are defined by occupation-level Software or Robot Exposure such that each quartile covers 1/4 of all workers in the 1980 sample, going from least to most exposed. Quartile assignments carry over to each year, regardless of employment share. Data: Weighted data from American Community Survey (1980-2010), restricted to employed individuals ages 16-64 living in a metropolitan area (AMA) (Ruggles et al., 2021). Exposure measurement from Webb (2019).

Figure A4. Change in Share of Male and Female Labor Force over Technology Exposure - Occupation-Level

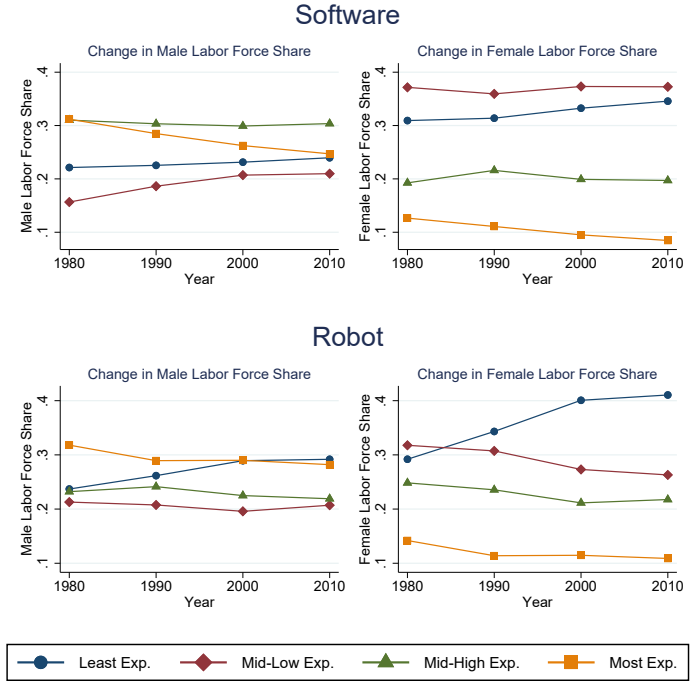

NOTE.— Above left: Change in the share of the male labor force employed in each software exposure quartile (1980 - 2010). Above right: Change in the share of the female labor force employed in each software exposure quartile (1980 - 2010). Below left: Change in the share of the male labor force employed in each robot exposure quartile (1980 - 2010). Below right: Change in the share of the female labor force employed in each robot exposure quartile (1980 - 2010). Quartiles are created separately for each technology category and defined by occupation-level Software or Robot Exposure such that each quartile covers 1/4 of all workers in the 1980 sample, going from least to most exposed. Quartile assignments carry over to each year, regardless of employment share. Data: Weighted data from American Community Survey (1980-2010), restricted to employed individuals ages 16-64 living in a metropolitan area (AMA) (Ruggles et al., 2021). Exposure measurement from Webb (2019).

Figure A5. Histograms of City-Level Technology Exposure

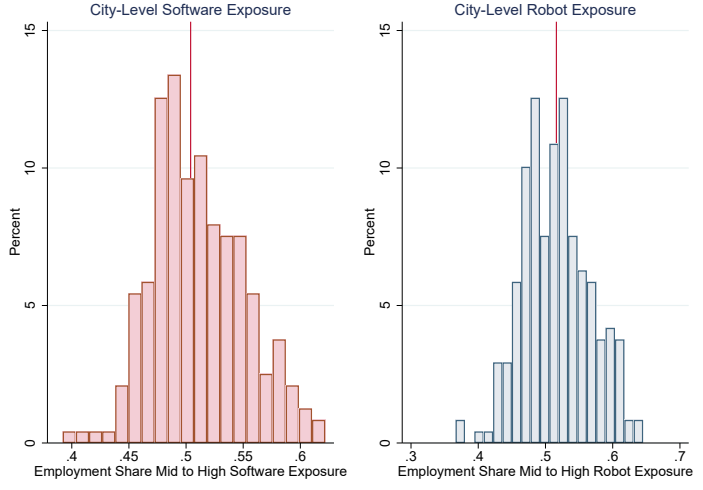

NOTE.— Histogram showing the distribution of city-level employment shares of the above median software and robot technology exposed occupations. Median value marked by vertical line. Data: Weighted data from American Community Survey (1980), restricted to employed individuals ages 16-64 living in a metropolitan area (AMA) (Ruggles et al., 2021). Each observation is a city that exists in the data in 1980 and 2010. Exposure measurement from Webb (2019).

Table A4. *Female Share and Occupation-Level Gender Wage Gap and Technology Exposure*

	(1)	(2)	(3)	(4)	(5)	(6)
	Δ Labor Force Share 1980-2010					
	Female			Male		
Software Exposure until 2018	-0.560***		0.061	-0.583***		-0.419
	(0.208)		(0.254)	(0.185)		(0.265)
Robot Exposure until 2018		-0.874***	-0.009***		-0.499***	-0.002
		(0.192)	(0.002)		(0.169)	(0.002)
Constant	0.343***	0.468***	0.456***	0.420***	0.371***	0.450***
	(0.115)	(0.110)	(0.123)	(0.098)	(0.100)	(0.107)
Observations	324	324	324	324	324	324
R^2	0.05	0.12	0.12	0.06	0.05	0.07
Mean Dependent Variable	0.089	0.089	0.089	0.155	0.155	0.155

	(1)	(2)	(3)	(4)	(5)	(6)
	Δ Female Share 1980-2010			Δ Occ-Level GWG 1980-2010		
Software Exposure until 2018	0.026		0.302**	-0.065*		-0.104**
	(0.111)		(0.135)	(0.033)		(0.049)
Robot Exposure until 2018		-0.221**	-0.004***		-0.006	0.001
		(0.104)	(0.001)		(0.037)	(0.001)
Constant	0.155**	0.262***	0.205***	-0.128***	-0.155***	-0.135***
	(0.062)	(0.055)	(0.061)	(0.021)	(0.025)	(0.023)
Observations	323	323	323	322	322	322
R^2	0.00	0.03	0.07	0.02	0.00	0.03
Mean Dependent Variable	0.167	0.167	0.167	-0.157	-0.157	-0.157

NOTE.— Dependent variables are the nominal difference in an occupation's share of female or male labor force employed in that occupation, female share of workers or the within-occupation gender wage gap from 1980 to 2010. Each observation is an occupation in the final sample. Software and Robot Exposure is defined as the exposure percentile, divided by 100 for legibility reasons. Observations are weighted by the mean labor supply size between 1980 and 2010. Data: Weighted data from American Community Survey (1980-2010), restricted to employed individuals ages 16-64 living in a metropolitan area (AMA) (Ruggles et al., 2021). Exposure measurement from Webb (2019).
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A5. *Change in Labor Demand - Software and Robot Exposure - Alternate End-Years*

Δ Empl. Share	(1)	(2)	(3)	(4)	(5)	(6)
			End Year			
	1990	2000	2019	1990	2000	2019
Software Exposure until 2018	-0.204**	-0.342**	-0.326			
	(0.090)	(0.158)	(0.218)			
Robot Exposure until 2018				-0.313***	-0.503***	-0.501***
				(0.084)	(0.152)	(0.193)
Constant	0.098*	0.185**	0.223**	0.143***	0.247***	0.292***
	(0.055)	(0.081)	(0.113)	(0.052)	(0.091)	(0.107)
Observations	324	323	324	324	323	324
R^2	0.03	0.03	0.02	0.08	0.07	0.04
Mean Dependent Variable	0.001	0.025	0.072	0.001	0.025	0.072
Δ Ln(Weekly Wage)	(1)	(2)	(3)	(4)	(5)	(6)
			End Year			
	1990	2000	2019	1990	2000	2019
Software Exposure until 2018	-0.116***	-0.190***	-0.232**			
	(0.024)	(0.053)	(0.091)			
Robot Exposure until 2018				-0.164***	-0.254***	-0.391***
				(0.024)	(0.050)	(0.080)
Constant	0.094***	0.240***	0.335***	0.114***	0.264***	0.398***
	(0.014)	(0.030)	(0.044)	(0.015)	(0.030)	(0.043)
Observations	324	323	305	324	323	305
R^2	0.12	0.11	0.07	0.26	0.21	0.22
Mean Dependent Variable	0.040	0.152	0.228	0.040	0.152	0.228

NOTE.— Top Panel: Dependent variable is the long difference in employment shares from 1980 to different end-years from 1990 to 2019, calculated as arc percentage changes. Bottom Panel: Dependent variable is the long difference in Ln(Weekly Wage) from 1980 to different end-years from 1990 to 2019. Each observation is an occupation, defined by Dorn et al. (2009) occ1990dd. Software and Robot Exposure is measured at the occupation level and is defined as Software or Robot Exposure Percentile/100 for legibility reasons. All columns use observations weighted by occupation cell's labor supply, averaged between 1980 and and the end-year used for that column. Offshorability measures occupation-level eligibility of offshoring, created by Autor and Dorn (2013). Data: Weighted data from US Census & American Community Survey (1980-2019), restricted to employed inidividuals ages 16-64 living in a metropolitan area (AMA) (Ruggles et al., 2021). Wages are winsorized at the bottom 1%. Exposure measurement from Webb (2019). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A6. *Change in Labor Demand - Software and Robot Exposure Interaction*

Δ Employment Share	(1)	(2)	(3)	(4)
High Software Exp.	-0.275**		-0.213*	
	(0.108)		(0.117)	
High Robot Exp.		-0.251**	-0.175	
		(0.111)	(0.114)	
Low Robot & High Software Exp.				-0.143
				(0.181)
High Robot & Low Software Exp.				-0.100
				(0.148)
High Robot & High Software Exp.				-0.395***
				(0.129)
Constant	0.169**	0.155*	0.218**	0.197**
	(0.076)	(0.084)	(0.091)	(0.099)
Observations	330	330	330	330
R^2	0.05	0.04	0.06	0.07
Mean Dependent Variable	0.044	0.044	0.044	0.044

Δ Ln(Weekly Wages)	(1)	(2)	(3)	(4)
High Software Exp.	-0.128***		-0.063*	
	(0.044)		(0.034)	
High Robot Exp.		-0.206***	-0.183***	
		(0.033)	(0.032)	
Low Robot & High Software Exp.				-0.066
				(0.056)
High Robot & Low Software Exp.				-0.187***
				(0.037)
High Robot & High Software Exp.				-0.246***
				(0.042)
Constant	0.228***	0.260***	0.279***	0.280***
	(0.025)	(0.025)	(0.028)	(0.031)
Observations	328	328	328	328
R^2	0.09	0.24	0.26	0.26
Mean Dependent Variable	0.169	0.169	0.169	0.169

NOTE.— Top Panel: Dependent variable is the long difference in employment shares from 1980 to 2010, calculated as arc percentage changes. Bottom Panel: Dependent variable is the long difference in Ln(Weekly Wage) from 1980 to 2010. Each observation is an occupation, defined by Dorn et al. (2009) occ1990dd. High Software or Robot Exposure is measured at the occupation level and equal to 1 if the occupation is above median exposed. All columns use observations weighted by occupation cell's labor supply, averaged between 1980 and 2010. Data: Weighted data from US Census & American Community Survey (1980-2010), restricted to employed inidividuals ages 16-64 living in a metropolitan area (AMA) (Ruggles et al., 2021). Wages are winsorized at the bottom 1%. Exposure measurement from Webb (2019).
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A7. *Employment-to-Population Ratio and City-Level Technology Exposure*

Δ Empl.-to-Pop. 1980-2010	(1)	(2)	(3)
$Exp_{c,1980}^{Software}$	-0.183***		-0.204***
	(0.057)		(0.074)
$Exp_{c,1980}^{Robot}$		-0.112	0.046
		(0.072)	(0.086)
Constant	0.325***	0.297***	0.308***
	(0.046)	(0.053)	(0.051)
Observations	239	239	239
R^2	0.53	0.51	0.53
Mean Dependent Variable	0.002	0.002	0.002
Demographic Controls	✓	✓	✓
Δ Female Empl.-to-Pop. 1980-2010	(1)	(2)	(3)
$Exp_{c,1980}^{Software}$	-0.188**		-0.346***
	(0.078)		(0.106)
$Exp_{c,1980}^{Robot}$		0.073	0.340***
		(0.100)	(0.118)
Constant	0.558***	0.411***	0.430***
	(0.068)	(0.076)	(0.070)
Observations	239	239	239
R^2	0.57	0.56	0.59
Mean Dependent Variable	0.055	0.055	0.055
Demographic Controls	✓	✓	✓
Δ Male Empl.-to-Pop. 1980-2010	(1)	(2)	(3)
$Exp_{c,1980}^{Software}$	-0.155**		-0.049
	(0.062)		(0.065)
$Exp_{c,1980}^{Robot}$		-0.265***	-0.227***
		(0.077)	(0.085)
Constant	0.056	0.139***	0.142***
	(0.037)	(0.048)	(0.049)
Observations	239	239	239
R^2	0.30	0.33	0.33
Mean Dependent Variable	-0.056	-0.056	-0.056
Demographic Controls	✓	✓	✓

NOTE.— Top Panel: Dependent variable is the nominal difference in a city's employment-to-population ratio from 1980 to 2010. Middle Panel: Dependent variable is the nominal difference in a city's female employment-to-population ratio from 1980 to 2010. Bottom Panel: Dependent variable is the nominal difference in a city's male employment-to-population ratio from 1980 to 2010. Each observation is a city in the final sample that exists in the data from 1950 to 2010. Software and Robot Exposure is defined as the employment share of the above median exposed occupations within a city 1980. Observations are weighted by 1980 labor supply size. City-level demographics include ethnicity and educational background and are measured in 1980. Data: Weighted data from American Community Survey (1980), restricted to employed inidividuals ages 16-64 living in a metropolitan area (AMA) (Ruggles et al., 2021). Wages are winsorized at the bottom 1%. 1980 wages are inflated to 2010 levels. Exposure measurement from Webb (2019). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A8. *Share of College Educated Labor Supply and City-Level Technology Exposure*

Δ College Share 1980-2010	(1)	(2)	(3)
$Exp_{c,1980}^{Software}$	0.042		0.357
	(0.214)		(0.236)
$Exp_{c,1980}^{Robot}$		-0.400**	-0.675***
		(0.159)	(0.166)
Constant	0.132	0.406***	0.387***
	(0.140)	(0.121)	(0.125)
Observations	239	239	239
R^2	0.29	0.32	0.34
Mean Dependent Variable	0.158	0.158	0.158
Demographic Controls	✓	✓	✓

Δ Female College Share 1980-2010	(1)	(2)	(3)
$Exp_{c,1980}^{Software}$	0.126		0.480**
	(0.210)		(0.228)
$Exp_{c,1980}^{Robot}$		-0.388**	-0.758***
		(0.154)	(0.159)
Constant	0.085	0.396***	0.370***
	(0.138)	(0.120)	(0.122)
Observations	239	239	239
R^2	0.35	0.37	0.41
Mean Dependent Variable	0.207	0.207	0.207
Demographic Controls	✓	✓	✓

NOTE.— Top Panel: Dependent variable is the level difference in a city's share of college educated from 1980 to 2010. Bottom Panel: Dependent variable is the level difference in a city's college share among female workers only from 1980 to 2010. Each observation is a city in the final sample that exists in the data from 1950 to 2010. Software and Robot Exposure is defined as the employment share of the above median exposed occupations within a city 1980. Observations are weighted by 1980 labor supply size. City-level demographics include ethnicity and educational background and are measured in 1980. Data: Weighted data from American Community Survey (1980), restricted to employed inidividuals ages 16-64 living in a metropolitan area (AMA) (Ruggles et al., 2021). Wages are winsorized at the bottom 1%. 1980 wages are inflated to 2010 levels. Exposure measurement from Webb (2019). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A9. *Gender Wage Gap and City-Level Software Technology Exposure (1980 - 2010): Added Controls*

Δ Residual GWG	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$Exp_{c,1980}^{Software}$	-0.324***	-0.357**	-0.361**	-0.379***	-0.457***	-0.514***	-0.439***	-0.485***
	(0.118)	(0.154)	(0.177)	(0.131)	(0.090)	(0.118)	(0.150)	(0.112)
Share Manufacturing	-0.080*	-0.026	-0.026	-0.079*				
	(0.042)	(0.059)	(0.059)	(0.042)				
Share Personal Services	0.020	1.268*	1.273*	0.022				
	(0.115)	(0.675)	(0.700)	(0.121)				
Δ Res. $GWG_{1950-1970}$		0.000	0.000			0.008	0.007	
		(0.042)	(0.042)			(0.040)	(0.040)	
$Exp_{c,1980}^{Robot}$			0.010	0.116			-0.152	0.052
			(0.237)	(0.143)			(0.242)	(0.158)
Share Offshorable Occ.					-0.139	-0.130	-0.174	-0.120
					(0.115)	(0.159)	(0.155)	(0.127)
Constant	-0.236***	0.434***	0.436***	-0.279***	-0.131	-0.127	-0.048	-0.157
	(0.073)	(0.083)	(0.093)	(0.092)	(0.081)	(0.101)	(0.158)	(0.117)
Observations	239	106	106	239	239	106	106	239
R^2	0.46	0.59	0.59	0.46	0.45	0.58	0.58	0.45
Mean Dep. Variable	-0.171	-0.171	-0.171	-0.171	-0.171	-0.171	-0.171	-0.171
Demographic Controls	✓	✓	✓	✓	✓	✓	✓	✓

NOTE.— Dependent variable is the nominal difference in a city's gender wage gap from 1980 to 2010. The residual gender wage gap uses residualized wages based on age, ethnicity, and educational background. Each observation is a city in the final sample that exists in the data from 1950 to 2010. Software exposure is defined as the employment share of the above median exposed occupations within a city in 1980. Observations are weighted by the size of the 1980 labor supply. City-level demographics include ethnicity and educational background and are measured in 1980. Data: Weighted data from the American Community Survey (1980), restricted to employed individuals ages 16 to 64 years living in a metropolitan area (AMA) (Ruggles et al., 2021). Wages are winsorized at the bottom 1%. 1980 wages are inflated to 2010 levels. Exposure measurement from Webb (2019). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A10. *Gender Wage Gap and City-Level Robot Technology Exposure (1980 - 2010): Added Controls*

Δ Residual GWG	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$Exp_{c,1980}^{Robot}$	-0.075	-0.184	0.010	0.116	-0.369***	-0.577***	-0.152	0.052
	(0.136)	(0.217)	(0.237)	(0.143)	(0.125)	(0.184)	(0.242)	(0.158)
Share Manufacturing	-0.145***	-0.089*	-0.026	-0.079*				
	(0.036)	(0.052)	(0.059)	(0.042)				
Share Personal Services	0.029	1.354*	1.273*	0.022				
	(0.122)	(0.694)	(0.700)	(0.121)				
Δ Res. $GWG_{1950-1970}$		-0.013	0.000			-0.010	0.007	
		(0.041)	(0.042)			(0.040)	(0.040)	
$Exp_{c,1980}^{Software}$			-0.361**	-0.379***			-0.439***	-0.485***
			(0.177)	(0.131)			(0.150)	(0.112)
Share Offshorable Occ.					-0.354***	-0.381**	-0.174	-0.120
					(0.118)	(0.159)	(0.155)	(0.127)
Constant	-0.338***	0.353***	0.436***	-0.279***	-0.048	0.088	-0.048	-0.157
	(0.092)	(0.088)	(0.093)	(0.092)	(0.110)	(0.142)	(0.158)	(0.117)
Observations	239	106	106	239	239	106	106	239
R^2	0.44	0.58	0.59	0.46	0.42	0.56	0.58	0.45
Mean Dep. Variable	-0.171	-0.171	-0.171	-0.171	-0.171	-0.171	-0.171	-0.171
Demographic Controls	✓	✓	✓	✓	✓	✓	✓	✓

NOTE.— Dependent variable is the nominal difference in a city's gender wage gap from 1980 to 2010. The residual gender wage gap uses residualized wages based on age, ethnicity, and educational background. Each observation is a city in the final sample that exists in the data from 1950 to 2010. Robot exposure is defined as the employment share of the above median exposed occupations within a city in 1980. Observations are weighted by the size of the 1980 labor supply. City-level demographics include ethnicity and educational background and are measured in 1980. Data: Weighted data from the American Community Survey (1980), restricted to employed individuals ages 16 to 64 years living in a metropolitan area (AMA) (Ruggles et al., 2021). Wages are winsorized at the bottom 1%. 1980 wages are inflated to 2010 levels. Exposure measurement from Webb (2019). * p $<$ 0.10, ** p $<$ 0.05, *** $p < 0.01$

Table A11. *Pre-Trend Analysis: Gender Wage Gap and City-Level Technology Exposure (1950 - 1970)*

	Δ GWG		Δ Residual GWG	
	(1)	(2)	(3)	(4)
Software Exposure				
$Exp_{c,1980}^{Software}$	0.069	0.796***	0.628***	0.605*
	(0.325)	(0.303)	(0.236)	(0.350)
$Exp_{c,1980}^{Robot}$				0.054
				(0.628)
Constant	0.059	-0.677***	-0.467***	-0.491
	(0.164)	(0.208)	(0.172)	(0.316)
Observations	106	106	106	106
R^2	0.00	0.28	0.11	0.11
Mean Dependent Variable	0.093	0.093	0.071	0.071
Demographic Controls		✓	✓	✓
	Δ GWG		Δ Residual GWG	
	(1)	(2)	(3)	(4)
Robot Exposure				
$Exp_{c,1980}^{Robot}$	-0.464*	0.786	0.615	0.054
	(0.237)	(0.543)	(0.455)	(0.628)
$Exp_{c,1980}^{Software}$				0.605*
				(0.350)
Constant	0.315***	-0.780**	-0.545*	-0.491
	(0.118)	(0.378)	(0.313)	(0.316)
Observations	106	106	106	106
R^2	0.05	0.25	0.09	0.11
Mean Dependent Variable	0.093	0.093	0.071	0.071
Demographic Controls		✓	✓	✓

NOTE.— Dependent variable is the nominal difference in a city's gender wage gap from 1980 to 2010. Residual gender wage gap uses wages residualized on age, ethnicity and educational background. Each observation is a city in the final sample that exists in the data from 1950 to 2010. Software and Robot Exposure is defined as the employment share of the above median exposed occupations within a city 1980. Observations are weighted by 1980 labor supply size. City-level demographics include ethnicity and educational background and are measured in 1980. Data: Weighted data from American Community Survey (1980), restricted to employed inidividuals ages 16-64 living in a metropolitan area (AMA) (Ruggles et al., 2021). Wages are winsorized at the bottom 1%. 1980 wages are inflated to 2010 levels. Exposure measurement from Webb (2019).
* p $<$ 0.10, ** p $<$ 0.05, *** $p < 0.01$

Table A12. *Robustness Check: Gender Wage Gap and City-Level Technology Exposure (1980 - 2010) - Holding number of cities fixed*

	Δ GWG		Δ Residual GWG			
	(1)	(2)	(3)	(4)	(5)	(6)
Software Exposure						
$Exp^{Software}_{c,1980}$	-0.760***	-0.629***	-0.537***	-0.543***	-0.519***	-0.513***
	(0.183)	(0.167)	(0.116)	(0.117)	(0.150)	(0.148)
Δ Res. GWG 1950-1970				0.010	0.011	
				(0.041)	(0.041)	
$Exp^{Robot}_{c,1980}$					-0.055	-0.055
					(0.243)	(0.242)
Constant	0.187**	-0.084	-0.172**	-0.167**	-0.143	-0.148
	(0.092)	(0.125)	(0.084)	(0.084)	(0.139)	(0.139)
Observations	106	106	106	106	106	106
R^2	0.22	0.42	0.58	0.58	0.58	0.58
Mean Dependent Variable	-0.191	-0.191	-0.171	-0.171	-0.171	-0.171
Demographic Controls		✓	✓	✓	✓	✓

	Δ GWG		Δ Residual GWG			
	(1)	(2)	(3)	(4)	(5)	(6)
Robot Exposure						
$Exp^{Robot}_{c,1980}$	-0.540***	-0.810***	-0.531***	-0.526***	-0.055	-0.055
	(0.160)	(0.236)	(0.184)	(0.188)	(0.243)	(0.242)
Δ Res. GWG 1950-1970				-0.008	0.011	
				(0.043)	(0.041)	
$Exp^{Software}_{c,1980}$					-0.519***	-0.513***
					(0.150)	(0.148)
Constant	0.068	0.121	-0.102	-0.106	-0.143	-0.148
	(0.080)	(0.180)	(0.129)	(0.132)	(0.139)	(0.139)
Observations	106	106	106	106	106	106
R^2	0.17	0.41	0.53	0.53	0.58	0.58
Mean Dependent Variable	-0.191	-0.191	-0.171	-0.171	-0.171	-0.171
Demographic Controls		✓	✓	✓	✓	✓

NOTE.— Dependent variable is the nominal difference in a city's gender wage gap from 1980 to 2010. Residual gender wage gap uses wages residualized on age, ethnicity and educational background. Each observation is a city in the final sample that exists in the data from 1950 to 2010. Software and Robot Exposure is defined as the employment share of the above median exposed occupations within a city 1980. Observations are weighted by 1980 labor supply size. City-level demographics include ethnicity and educational background and are measured in 1980. Data: Weighted data from American Community Survey (1980), restricted to employed inidividuals ages 16-64 living in a metropolitan area (AMA) (Ruggles et al., 2021). Wages are winsorized at the bottom 1%. 1980 wages are inflated to 2010 levels. Exposure measurement from Webb (2019).
* p $<$ 0.10, ** p $<$ 0.05, *** $p < 0.01$

Table A13. *Robustness Check: Gender Wage Gap and City-Level Technology Exposure (1980-2010) - Only Most Exposed Occupations*

	Δ GWG		Δ Residual GWG				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Software Exposure							
$Exp^{Software}_{c,1980}[Qrt.4]$	-0.394**	-0.373**	-0.315***	-0.500***	-0.503***	-0.547**	-0.320**
	(0.186)	(0.154)	(0.108)	(0.173)	(0.166)	(0.235)	(0.128)
Δ Res. GWG 1950-1970				-0.003		-0.002	
				(0.044)		(0.045)	
$Exp^{Robot}_{c,1980}[Qrt.4]$						0.070	0.007
						(0.280)	(0.179)
Constant	-0.099**	-0.294***	-0.358***	-0.317***	-0.315***	-0.332***	-0.360***
	(0.048)	(0.078)	(0.053)	(0.076)	(0.073)	(0.104)	(0.064)
Observations	239	239	239	106	106	106	239
R^2	0.05	0.28	0.40	0.53	0.53	0.53	0.40
Mean Dependent Variable	-0.192	-0.192	-0.171	-0.171	-0.171	-0.171	-0.171
Demographic Controls		✓	✓	✓	✓	✓	✓

	Δ GWG		Δ Residual GWG				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Robot Exposure							
$Exp^{Robot}_{c,1980}[Qrt.4]$	-0.338**	-0.388*	-0.259*	-0.392*	-0.395*	0.070	0.007
	(0.146)	(0.210)	(0.153)	(0.222)	(0.220)	(0.280)	(0.179)
Δ Res. GWG 1950-1970				-0.021		-0.002	
				(0.043)		(0.045)	
$Exp^{Software}_{c,1980}[Qrt.4]$						-0.547**	-0.320**
						(0.235)	(0.128)
Constant	-0.109***	-0.248**	-0.345***	-0.293***	-0.289***	-0.332***	-0.360***
	(0.039)	(0.100)	(0.066)	(0.104)	(0.103)	(0.104)	(0.064)
Observations	239	239	239	106	106	106	239
R^2	0.05	0.28	0.39	0.52	0.52	0.53	0.40
Mean Dependent Variable	-0.192	-0.192	-0.171	-0.171	-0.171	-0.171	-0.171
Demographic Controls		✓	✓	✓	✓	✓	✓

NOTE.— Dependent variable is the nominal difference in a city's gender wage gap from 1980 to 2010. Each observation is a city in the final sample that exists in the data from 1950 to 2010. Software and Robot Exposure is defined as the employment share of the most exposed quartile of occupations within a city 1980. Observations are weighted by 1980 labor supply size. City-level demographics include ethnicity and educational background and are measured in 1980. Data: Weighted data from American Community Survey (1980), restricted to employed inidividuals ages 16-64 living in a metropolitan area (AMA) (Ruggles et al., 2021). Wages are winsorized at the bottom 1%. 1980 wages are inflated to 2010 levels. Exposure measurement from Webb (2019).
* p $<$ 0.10, ** p $<$ 0.05, *** $p < 0.01$

Table A14. *Robustness Check: Gender Wage Gap and City-Level Technology Exposure (1980-2010) - Continuous Measurement*

	Δ GWG		Δ Residual GWG				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Software Exposure							
$Exp^{Software}_{c,1980}[Cont.]$	-0.008***	-0.008***	-0.008***	-0.010***	-0.010***	-0.008*	-0.008***
	(0.003)	(0.002)	(0.001)	(0.002)	(0.002)	(0.004)	(0.002)
Δ Res. GWG 1950-1970				0.013		0.011	
				(0.041)		(0.041)	
$Exp^{Robot}_{c,1980}[Cont.]$						-0.003	0.000
						(0.005)	(0.003)
Constant	0.180	0.585***	0.542***	0.041	0.037	0.129	0.544***
	(0.135)	(0.130)	(0.088)	(0.129)	(0.129)	(0.177)	(0.088)
Observations	239	239	239	106	106	106	239
R^2	0.09	0.31	0.44	0.58	0.57	0.58	0.44
Mean Dependent Variable	-0.192	-0.192	-0.171	-0.171	-0.171	-0.171	-0.171
Demographic Controls		✓	✓	✓	✓	✓	✓

	Δ GWG		Δ Residual GWG				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Robot Exposure							
$Exp^{Robot}_{c,1980}[Cont.]$	-0.006***	-0.010***	-0.008***	-0.011***	-0.011***	-0.003	0.000
	(0.002)	(0.003)	(0.002)	(0.003)	(0.003)	(0.005)	(0.003)
Δ Res. GWG 1950-1970				-0.003		0.011	
				(0.041)		(0.041)	
$Exp^{Software}_{c,1980}[Cont.]$						-0.008*	-0.008***
						(0.004)	(0.002)
Constant	0.098	0.459***	0.397***	0.833***	0.835***	0.129	0.544***
	(0.096)	(0.115)	(0.081)	(0.200)	(0.197)	(0.177)	(0.088)
Observations	239	239	239	106	106	106	239
R^2	0.09	0.30	0.42	0.56	0.56	0.58	0.44
Mean Dependent Variable	-0.192	-0.192	-0.171	-0.171	-0.171	-0.171	-0.171
Demographic Controls		✓	✓	✓	✓	✓	✓

NOTE.— Dependent variable is the nominal difference in a city's gender wage gap from 1980 to 2010. Each observation is a city in the final sample that exists in the data from 1950 to 2010. Software and Robot Exposure for each city is defined as the employment weighted percentile score in 1980. Observations are weighted by 1980 labor supply size. City-level demographics include ethnicity and educational background and are measured in 1980. Data: Weighted data from American Community Survey (1980), restricted to employed inidividuals ages 16-64 living in a metropolitan area (AMA) (Ruggles et al., 2021). Wages are winsorized at the bottom 1%. 1980 wages are inflated to 2010 levels. Exposure measurement from Webb (2019). * p < 0.10, ** p < 0.05, *** $p < 0.01$

Table A15. *Robustness Check: Gender Wage Gap and City-Level Technology Exposure - Alternate End-Years 1990 - 2019*

End-Year	1990		2000		2019	
Δ Residual GWG	(1)	(2)	(3)	(4)	(5)	(6)
Software Exposure						
$Exp_{c,1980}^{Software}$	-0.012	-0.124	-0.103	-0.232*	-0.452**	-0.402*
	(0.110)	(0.125)	(0.099)	(0.121)	(0.185)	(0.213)
$Exp_{c,1980}^{Robot}$		0.242*		0.276**		-0.120
		(0.129)		(0.124)		(0.213)
Constant	-0.174**	-0.267***	-0.311***	-0.415***	-0.256*	-0.208
	(0.080)	(0.089)	(0.069)	(0.072)	(0.132)	(0.152)
Observations	219	219	239	239	174	174
R^2	0.21	0.22	0.28	0.30	0.47	0.47
Mean Dependent Variable	-0.108	-0.108	-0.152	-0.152	-0.162	-0.162
Demographic Controls	✓	✓	✓	✓	✓	✓

End-Year	1990		2000		2019	
Δ Residual GWG	(1)	(2)	(3)	(4)	(5)	(6)
Robot Exposure						
$Exp_{c,1980}^{Robot}$	0.146	0.242*	0.098	0.276**	-0.447**	-0.120
	(0.124)	(0.129)	(0.104)	(0.124)	(0.203)	(0.213)
$Exp_{c,1980}^{Software}$		-0.124		-0.232*		-0.402*
		(0.125)		(0.121)		(0.213)
Constant	-0.273***	-0.267***	-0.428***	-0.415***	-0.213	-0.208
	(0.090)	(0.089)	(0.072)	(0.072)	(0.149)	(0.152)
Observations	219	219	239	239	174	174
R^2	0.22	0.22	0.28	0.30	0.44	0.47
Mean Dependent Variable	-0.108	-0.108	-0.152	-0.152	-0.162	-0.162
Demographic Controls	✓	✓	✓	✓	✓	✓

NOTE.— Dependent variable is the nominal difference in a city's gender wage gap from 1980 to 2010. Each observation is a city in the final sample that exists in the data from 1950 to 2010. Software and Robot Exposure is defined as the employment share of the above median exposed occupations within a city 1980. Observations are weighted by 1980 labor supply size. City-level demographics include ethnicity and educational background and are measured in 1980. Data: Weighted data from American Community Survey (1980), restricted to employed inidividuals ages 16-64 living in a metropolitan area (AMA) (Ruggles et al., 2021). Wages are winsorized at the bottom 1%. 1980 wages are inflated to 2010 levels. Exposure measurement from Webb (2019).
* p $<$ 0.10, ** p $<$ 0.05, *** $p < 0.01$

Table A16. *Ln(Weekly Wages) and Occupation-Level AI Exposure Percentile (2019)*

Ln(Weekly Wage)$_{i,2019}$	(1)	(2)	(3)	(4)	(5)
AI Exposure Percentile	0.819***	0.788***	0.761***	0.312***	0.310***
	(0.004)	(0.004)	(0.005)	(0.004)	(0.004)
Female		-0.113***	-0.141***	-0.163***	-0.160***
		(0.002)	(0.004)	(0.002)	(0.002)
Female x AI Exp.			0.057***		
			(0.007)		
Average Age				0.014***	0.014***
				(0.000)	(0.000)
Share Black				-0.167***	-0.165***
				(0.003)	(0.003)
Share Hispanic				-0.108***	-0.133***
				(0.002)	(0.003)
Share High School				0.142***	0.150***
				(0.004)	(0.004)
Share Some College				0.264***	0.271***
				(0.004)	(0.004)
Share College				0.678***	0.662***
				(0.004)	(0.004)
Observations	991,931	991,931	991,931	991,931	991,931
R^2	0.09	0.09	0.09	0.38	0.40
Weighted Obs.	118,931,944	118,931,944	118,931,944	118,931,944	118,931,944
Industry Dummies				✓	✓
AMA Dummies					✓

NOTE.— Dependent variable is Ln(Weekly Wage). AI Exposure is defined as AI Exposure Percentile/100 for legibility reasons. Industry dummies include all industries as defined by 'ind1990' in ACS. AMA dummies include all metropolitan areas that exist in the data from 1980 to 2010. The reference variable for ethnicity is White. The reference variable for Educational Background is 'Less than High School'. Data: Weighted data from American Community Survey (ACS, 2000), restricted to employed inidividuals ages 16-64 living in a metropolitan area (AMA) (Ruggles et al., 2021). Wages are winsorized at the bottom 1%. Exposure measurement from Webb (2019).
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Essay III. Are Economics Students Biased against Female Teachers? Evidence from a Randomized, Double-Blind Natural Field Experiment

Co-authored with Ola Andersson, Niklas Bengtsson, and Per Engström

Acknowledgements: All at the Department of Economics, Uppsala University. Corresponding author per.engstrom@nek.uu.se. We would like to extend our gratitude to Julia Strandberg for outstanding administrative support and to Elin Sundberg and Anton Sundberg for letting us use their names. We also thank Oskar Nordström Skans, Anna Dreber, Jan Sauermann, Magnus Johannesson and seminar participants at the Department of Economics at Uppsala University and Stockholm Behavioral Network workshop at SSE for their valuable comments.

1 Introduction

Student evaluations of teaching (SETs) are widely used to improve class instruction and organize higher education (Kulik, 2001; Benton and Cashin, 2013). Job market candidates and untenured professors routinely include their teacher ratings when applying for academic jobs or promotions (Moore and Trahan, 1998; Keng, 2018). However, recent studies indicate that SETs are biased against female teachers and may disadvantage women's academic careers (Boring and Ottoboni, 2016; Heffernan, 2021; Ceci et al., 2023).

Most existing studies on gender bias in SETs use observational data, an approach that potentially suffers from both selection and observability issues. Although some studies employ clever designs to reduce the problem of student-teacher selection, researchers cannot observe and control for all relevant differences between female and male instructors in the classroom. This research challenge is not specific to student evaluations, but to all observational studies of gender bias. For this reason, correspondence experiments have become the gold standard in studies of discriminatory behavior (Bertrand and Duflo, 2017). However, such experiments are difficult to set up in realistic learning environments where students and teachers have repeated interactions.

In this study, we estimate gender bias in SETs using a randomized, double-blind natural field experiment. Specifically, we created an experimental learning environment in which students believed they interacted with either a male or a female teacher – while the actual gender was orthogonal to the stated gender. The opportunity to implement the experiment arose naturally in the wake of the Covid-19 pandemic when part of the teaching had moved online. Students were instructed to direct questions to "email mentors" instead of on-campus teaching assistants. The regular course instructors provided replies, but following the research methods in correspondence studies, we randomly signed the answers to each student request (765 emails) with either a female or a male name. Importantly, the course instructors did not know whether they acted as female or male in their correspondence; the experiment was thus "double-blind", which is a novel contribution to the literature. After the course, the students rated the performance of the female and male email mentors separately.[1]

The results show no evidence of student bias against the female teacher in any considered dimension (the dimensions are 'Helpful,' 'Knowledgeable,' 'Fast,' and an 'Overall' rating). For the overall rating, the point estimate is 0.018 at a baseline rating of 4.6 out of 5 (0.027 of a standard deviation), and the 95%-level confidence interval is $\pm$ 0.094. The tight confidence bounds imply that the statistical power is enough to reject effect sizes found in the previous literature (Ceci et al., 2023). The zero estimates do not mask differential effects; the estimates are the same across female and male students (preregis-

[1]The study is preregistered at: https://osf.io/7rk5g and vetted by The Swedish Ethical Review Authority (Dnr 2021-05914-01).

tered subgroups). In a supplementary analysis (not preregistered), we find no evidence that female teachers face a more demanding teaching environment. The perceived female mentor did not trigger more follow-up questions or less polite email feedback (the verbal feedback from students ranged from neutral to positive).

Our findings contrast previous research. An important contribution is MacNell et al. (2015), an audit study where online teachers performed as either male or female. The authors report a relatively strong bias against female teachers. Similar to our study, their experiment randomized the perceived genders of instructors, but in their study, the actual instructors knew whether they performed as male or female. The convention in randomized trials is that the person administrating the treatment should be blind to treatment status to avoid (consciously or subconsciously) influencing the outcome.[2] Arguably, the concern of experimenter demand effects is particularly relevant in a setting where the treatment is performative (Bertrand and Duflo, 2017). Our prior, outlined in the preexperimental plan, is that knowledge about which gender the instructors are performing as risks affecting the style of instruction. By setting up an administrative intermediary between the instructors and the students, we ensured that instructors never knew their "own" gender. Furthermore, we logged the correspondence between students and mentors, allowing us to confirm through balance tests that the emails sent by the instructors do not differ between the two gendered mentors.

Previous research using observational data has found evidence of student bias against female teachers. Two of these studies resolve the student selection issue using allocation rules that match students to teachers in a quasi-random way. Using online evaluations from over 20,000 students, Boring (2017) finds that ratings are biased against female teachers and that male students drive this bias. Mengel et al. (2019) exploit a similar allocation rule at Maastricht University in the Netherlands and find that female instructors receive lower scores from both men and women (although senior female instructors are rated more favorably by female students). Although these studies control for selection, they cannot completely control for confounders (factors observed by the students but not by the researchers), which may bias the results.

Our paper is also related to the literature on gender bias in academic environments more generally (Moss-Racusin et al., 2012; Cheryan et al., 2017; Nielsen, 2016; Weisshaar, 2017). Field experiments on how students rate female-produced teaching materials have rendered mixed results (Arbuckle and Williams, 2003; Özgümüs et al., 2020; Binderkrantz et al., 2022). Regard-

[2]Quasi-experimental studies may also suffer from experimenter demand effects. Khazan et al. (2019) directed students to either a putative female or putative male online TA (although it was the same TA). They find data patterns supportive of strong bias against female teachers but lack the statistical power to reject zero effects. Wong and Bouchard (2021) alternated the perceived gender of a teaching assistant across academic semesters and find mixed evidence. None of these studies are double-blind, and the treatment assignment was nonrandom.

ing discipline-specific learning environments, there is some evidence that economics is especially plagued by sexism (see, e.g., Paredes et al. 2023). There can be many reasons for this, but a proposed factor is that economics has long been male-dominated and continues to be so to this date (Kahn, 1993; Kahn, 1995; Lundberg and Stearns, 2019). Aragón et al. (2023) find that course evaluations are biased against the underrepresented gender in an artefactual field experiment. A large literature documents a gendered role-model effect in higher education (Rothstein, 1995; Bettinger and Long, 2005; Hoffmann and Oreopoulos, 2009; Carrell et al., 2010; Paredes, 2014; Buckles, 2019). The Department of Economics at Uppsala University is also male-dominated, yet we find no SET bias against women in this setting.

With respect to the external validity of our findings, we note that our study is conducted online, with the name of the teacher being the only indicator of gender. It is unclear whether student bias against female teachers is more or less likely to manifest in text-based online environments. While it could be argued that actual contact with a female professor increases the "gender dosage" at a given level of prejudice, the contact hypothesis (Allport, 1954) suggests that such contacts may also decrease a student's prejudices.[3] In general, remote environments can be extremely toxic towards women (an example is online gaming; see Fox and Tang 2017). In higher education, there is evidence that remote environments increase negative comments about female teachers from male students (Tangalakis et al., 2022), and that semesters with remote teaching increase SET bias against women (Ayllón, 2022). Notably, the most widely-cited study in this field (MacNell et al., 2015) was conducted in a text-based environment similar to ours, and it reports the largest effect sizes in the literature (see the discussion in Ceci et al. 2023, as well as our effect size comparison in Section 5). According to Ceci et al. (2023), MacNell, Driscoll, and Hunt's study is "the most unequivocal experimental evidence of bias" to date.

This paper is organized as follows. The next section overviews related literature and motivates the research question and method. Section 3 describes the experiment. Section 4 describes the results, and Section 5 concludes.

2 Bias in SETs

Student evaluations of teaching have a long tradition in higher education; the first study on teacher ratings dates back to 1927 (Remmers and Brandenburg, 1927). Despite their long history, they remain controversial in academia. One criticism is that SETs are easily manipulated and do not necessarily reflect

[3]Related, Kreisman and Smith (2023) find that students with more statistically Black names have differential educational outcomes, but this relationship disappears when comparing siblings. They conclude that correspondence studies may be more likely to pick up biases since a teacher's physical meeting with a Black student may negate the adverse stereotype effects of having a distinctively Black name.

the quality of teaching, making them unsuitable for tenure decisions (Stroebe, 2020). Even when SETs have no direct role in tenure or hiring decisions, they indirectly affect career outcomes as they influence teacher-course matching (Kulik, 2001). Biased evaluations may also be harmful and discouraging in themselves (Bates, 2015).

Previous studies tend to show that women are rated less favorably. Hamermesh and Parker (2005) report that female instructors receive markedly lower ratings (although the study focuses on beauty, not gender). A recent study is Fan et al. (2019), which analyzes over 500,000 student evaluations and finds that female instructors receive lower average scores, even when students are asked to evaluate the *course*, and not the specific instructor. Another recent study on over 100,000 students finds no overall bias against females but a significant same-sex bias (Binderkrantz and Bisgaard, 2023). Analyses of RateMyProfessor-type of websites have produced mixed results. Although earlier studies found no differences across gender (Reid, 2010; Stuber et al., 2009), more recent studies report lower ratings for females (Rosen, 2018; Boehmer and Wood, 2017; Arceo-Gomez and Campos-Vazquez, 2019). A study of a student-driven Swedish website found that female teachers get lower ratings (Karlsson and Lundberg, 2012).

The gender bias found in previous studies is striking, especially since students have no personal stakes when answering anonymous course evaluations. Unlike correspondence experiments using job applications or research proposals, the stake-free situation means a detected bias is unlikely to reflect statistical discrimination.[4] The remaining explanation is that students have male-biased preferences. Relative to students in other disciplines, there is evidence that economics students are more likely to hold sexist beliefs when enrolling in higher education (Paredes et al., 2023). Sexism in anonymous online forums, such as *Economics Job Market Rumors*, is well-known and documented in Wu (2018). If such beliefs influence performance evaluations, anonymous SET ratings may differ between male and female teachers even when their performance is identical.

Observational studies on gender bias in teacher ratings need to take a stand on what parts of the variation in student satisfaction reflect a teacher's professional qualities and what parts reflect irrelevant, "biased" factors. Some previous research uses exam scores to control for the actual quality of teaching (e.g., Boring 2017 and Mengel et al. 2019), meaning that residual differences across the teachers' gender are interpreted as biases. However, whether grades and pass rates are better proxies for teacher quality than SET ratings is unclear. Teachers may compensate for low instructive quality by giving lenient courses or, if they do not control the exam, by employing a "teaching to the test"-

[4]Résumé experiments and audit studies have rendered mixed evidence of gender bias in the labor market (Neumark et al., 1996; Bertrand and Mullainathan, 2004; Riach and Rich, 2006; Petit, 2007; Carlsson and Eriksson, 2019; Carlsson, 2011; Ahmed et al., 2021).

pedagogy. There is some evidence that assistant professors routinely do so as the tenure decision approaches (Moore and Trahan, 1998; Keng, 2018).

Diverging SET ratings need not reflect biases, even at given (true) learning outcomes, as student evaluations of teaching can be informative about teaching externalities. For example, teachers may incur negative or positive externalities upon the study environment or the workload of other faculty members, teaching assistants, or course administrators.[5] These aspects may be observed by students but not by researchers. Another conceptual point is that departments may have an intrinsic motive to hire inspiring, entertaining, or provocative teachers even at given learning outcomes, if such teachers increase the consumption value of higher education.[6] Controlling for the consumption value of a specific instructor in observational studies is rarely possible.

An ideal control in observational studies would be an accurate measure of a professor's value added to students minus the externalities imposed on administration and colleagues. Such controls do not exist, partly because opinions diverge on what the values added and costs are (or should be). Observational studies may underestimate or overestimate student biases against female teachers if these abstract costs and benefits are not held constant. The benefit of correspondence experiments is that expected professional qualities need not be observed since they are identical across treatment and control groups by design.

3 The Experiment

The Department of Economics at Uppsala University gives an Introductory Economics course to about 250 students each semester. The course consists of lectures and tutorials where problem sets are solved. The department uses student mentors to facilitate the latter. These mentors are usually master students who help students solve exercises in smaller classes (the "mentoring" refers to guidance in solving problem sets only, it does not include career guidance). During the spring of 2021, when the Covid-19 pandemic restrictions prevented full campus participation, some mentoring took place online through email. This second-best solution was a success; the students were active in emailing questions and gave encouraging remarks in the course evaluations.

The success of the forced switch to text-based mentoring gave us the idea of reintroducing email mentoring even after Covid-19 was not restricting campus teaching anymore. So, in the spring of 2022, we let the students know they could email any questions about the course content to the generic email address *mailmentor@nek.uu.se* and that two of the department's Ph.D. students would answer their questions. In practice, the email mentors consisted of one

[5]For example, at Uppsala University, SETs routinely include questions about sexual harassment and equal opportunities.

[6]For the consumption value of higher education, see, e.g., Jacob et al. (2018)

single Ph.D.-student (Malin Backman) and the two principal teachers of the course (Per Engström, spring of 2022, and Ola Andersson, fall of 2022).

We hereafter refer to the persons answering the emails by the term "instructors." Only one administrator could access the generic email address at the department. The instructors could not access it; instead, the administrator forwarded the email queries from students. The instructor then emailed the response to the administrator, who randomly signed the answer with either Anton (male) or Elin (female), two common Swedish names. We used pairwise block randomization, meaning the mentor's name was randomized within every two emails sent from the mentor email. Answers to follow-up questions were also randomized.

The experimental protocol implied that a student could be exposed to a long email thread with answers from both male and female mentors. Appendix A.3 gives an example of a typical correspondence. A description of the example email is outlined here:

1. A student emails *mailmentor@nek.uu.se*. The email is forwarded to the instructors by the administrator in charge of the email account.
2. The instructor writes a response and sends it back to the administrator, who signs the email with either Anton or Elin, selecting the name from a preset list of mentor names in a random order. The preset list meant that the administrator had no discretion at this stage.
3. Sometimes, the student used their email client's reply button to send follow-up questions to the same address. The student's reply is forwarded by the administrator to the instructor.
4. The instructor answers the follow-up question and sends the answer to the administrator, who signs the email with a new random draw of Anton or Elin.

The perceived gender of the instructor is only identifiable through the name in the signature, following previous correspondence and audit studies similar to ours (Bertrand and Mullainathan, 2004). Both Elin and Anton are native Swedish names and send a clear signal of the sender's gender. According to Statistics Sweden, about 37,000 women are called (first name) Elin (only 14 men), and 28,500 men are called Anton (only 11 women). For added credibility, the names referred to actual doctoral students at the department. Both gave consent to use their names, and we instructed these two doctoral students to refer to the *mailmentor*-email address in case a student contacted them directly (this did not happen).

To measure the presence of a bias, we introduced three preregistered questions for each mentor's gender in the course evaluation. Specifically, students were asked to rate Elin and Anton on a five-point scale regarding helpfulness, knowledge, and response time. See Table A5 for the complete (translated) questions used in the evaluation. Although the dimensions reflect different gendered qualities, our prior was that responses would be biased against the

female mentor in all three dimensions.[7] With all dimensions unidirectional, we preregistered an 'Overall' outcome for maximum statistical power, equal to the unweighted average of the three dimensions. In cases where students only rate one or two dimensions, 'Overall' equals the average of the rated dimension(s). We randomized the order of the evaluation questions (half the students answered questions about Elin first and the other half about Anton first).

The setup meant the educational environment was perfectly natural from a student's perspective. However, the exact phrasing of the course evaluation questions created a slight deviation from our department's standards. The standard way of phrasing SET questions at Uppsala University is component/function, not person/human trait (e.g., "Were the lectures well-prepared?" rather than "Was Ola knowledgeable?").[8] We departed somewhat from these standards to align our study with existing literature. MacNell et al. (2015) use nine additional dimensions, but we could not include that many questions about mentors without compromising the natural frame of the course evaluation.

Since course evaluations are strictly anonymous in this context, we cannot link email interactions to specific answers in the course evaluation. This means we cannot test whether a male or female mail mentor affected participation in the course evaluation (which was voluntary). However, this participation effect is part of the mechanism we aim to estimate. In Section 5.3 we analyze whether the female mentor triggered other student responses, such as getting more follow-up questions or less encouraging direct feedback.

4 Estimation

Although the treatment is randomized, we preregistered several estimation methods to increase precision. Because the same student is asked to rate both the female and male mentors, a within-estimator, with student-fixed effects, is technically feasible. Using the same student's female-male rating difference solves the issue that very dissatisfied students may influence the estimation in a way that severely reduces precision.[9] Although the fixed-effects method is

[7] We primarily base this prior on previous empirical literature. However, Adams et al. (2021) discuss how gender stereotypes may, theoretically, interact with student ratings. They argue that female teachers are punished for showing too little helpfulness, not because they are less helpful but because they are expected to be more helpful. The stereotype that men are more technically skilled leads instead to confirmation bias, in that students interpret simple answers from male teachers as more advanced than they really are.

[8] Uppsala University guidelines (Avdelningen för kvalitetsutveckling, 2020) suggest that teacher names should be avoided as far as possible, out of the same concerns that motivate this research paper.

[9] Uttl and Violo (2021) indicate that a few such students may have affected the conclusions in MacNell et al. (2015).

more efficient at given response rates, it requires that students answer questions about both mentors.

We expected about 200 students would participate in the course evaluations when running the experiment over two semesters. However, we expected only a subset of students would interact with the mentors and, consequently, rate them. Furthermore, given that students could send multiple questions to the mail mentors and were treated with a random draw of the female and male mentor each time, it was at the experimental planning stage not clear how many students would be exposed to both genders.[10] In our preregistered power analysis, we estimated that for a fixed-effects approach to be feasible, at least 80 students had to rate both the female and male mentors (see Appendix B for the details of the power analysis).

Given these uncertainties, the preexperimental plan outlined three separate approaches: a comparison of mean responses (*mean*), regression-adjusted estimates using controls for the student's age, gender, and educational program (*covs*), and a fixed effects analysis (*fe*). Our preexperimental commitment was that *covs* would be our preferred method if the number of students rating both mentors were lower than 80; otherwise, we would employ a fixed-effects estimation.

To perform the estimation, we stack each mentor's rating (female, male) within each student j, meaning that the unit of observation i is a student-mentor pair. We cluster-adjust the standard errors on the student level. The regression specification is

$$y_{ij} = \gamma_0 + \gamma_1 \text{female}_{ij} + \delta \mathbf{X_j} + u_{ij}, \tag{1}$$

where y_{ij} is the student's rating of the mentors and the vector $\mathbf{X_j}$ either captures student fixed effects (`fe`) or controls for age, student gender, and educational program (`covs`). We run regression (1) on the full sample as well as across female and male students. Our parameter of interest is γ_1, which is negative if students are biased against females.

5 Results

5.1 Data and balancing tests

The access to email mentors proved popular, with 765 emails answered over two semesters from 203 unique student email addresses, around 42% of all students enrolled. Of these 203 students, 116 were female, and 87 were male. However, the phrasing of the student queries indicates that students

[10] On a conceptual note, it is possible that some students rate teachers they have never interacted with. Such behavior may add to the observed gender bias in SETs, and should thus be part of the estimation. On a technical note, the SET response rate could theoretically be larger than the number of email interactions.

often reached out to the mentor as a small group, so the total reach of the mentors is likely larger than 203 individuals. 153 unique email addresses had at least one interaction with the perceived female mentor, Elin, and 148 students had at least one interaction with the perceived male mentor, Anton; 98 students had at least one interaction with both iterations of the mentor.

Table 1. *Balance test of emails sent.*

	Male mentor	Female mentor	Difference
Female student	0.614	0.571	-0.0429
	(0.0249)	(0.0254)	(0.0355)
Length of answer	499.4	526.6	27.15
	(20.39)	(23.58)	(31.17)
Attachment	0.0392	0.0262	-0.0130
	(0.00993)	(0.00818)	(0.0129)
Response time (hours)	16.01	17.72	1.717
	(1.119)	(1.173)	(1.621)
Response time (log)	1.597	1.735	0.138
	(0.0948)	(0.0958)	(0.135)
Observations	383	382	765

NOTE.— 'Attachment' is a dummy indicating whether the mentor included an attachment in the response email. The answer length refers to the number of characters in the response email. The first two columns display standard deviations in parentheses, the third column displays the standard errors of the difference in means in parenthesis.

The distribution of emails is evenly spread across the terms (see Figure A1 in Appendix 6). In the end, "Anton" responded to 383 emails and "Elin" to 382. Table 1 shows that the treatment and control groups are precisely balanced and that the characteristics of the emails do not significantly differ between the mentors. Importantly, conditional upon interacting with any mentor, female students were not more or less likely to interact with the female mentor or vice versa. We also cannot reject that the average speed or length of the response was the same.

All in all, 253 students responded to the course evaluation, 136 female students (54 percent) and 117 male students (46 percent). Of these, 109 responded to at least one question about any mentor, and 78 students evaluated both the female and the male mentors. Table 2 presents descriptive statistics of students rating the mentors. The average rating of the email mentors turned out to be high (about 4.62 on a 5-point scale; see Table 2). As discussed in section 3, more experienced instructors responded to emails during the experiment (compared to the usual teaching assistants). This choice reflected two separate considerations. The first was that with increased quality, we hoped to maximize the popularity of the online mentors and hence participation in the experiment. The second consideration was ethical. To encourage participation, the course instructors informed students about the email mentors in

class (without saying it was part of an experiment). We could not risk that the experiment effectively lowered the standard of teaching.[11]

Table A1 in Appendix A replicates the summary statistics in Table 2 but separately for female and male students. Conditional upon filling out the course evaluation, 53 percent (72 out of 136) of the female students rated any mentor, but only 31 percent (37 of 117) of the male students did. This skewness reflects that male students were less likely to interact with online mentors. Conditional upon rating any mentor, female and male students were equally likely to rate both (0.71 vs 0.73).

Table 2. *Summary statistics: Course evaluation*

Ratings (Male mentor)	Mean	Std. dev.	Obs.
Helpful	4.644	0.754	90
Fast in replying	4.541	0.795	85
Knowledgeable	4.713	0.680	87
Overall (average)	4.635	0.679	90
Ratings (Female mentor)	Mean	Std. dev.	Obs.
Helpful	4.646	0.725	96
Fast in replying	4.489	0.826	94
Knowledgeable	4.663	0.724	95
Overall (average)	4.608	0.701	97
Student characteristics	Mean	Std. dev.	Obs.
Male student	0.34	0.476	109
<20 years old	0.04	0.189	109
20 years old	0.21	0.410	109
21 years old	0.25	0.434	109
22 years old	0.17	0.381	109
23 years old	0.10	0.303	109
24 years old	0.12	0.326	109
>24 years	0.11	0.314	109
Major: Business	0.53	0.501	109
Major: Political science	0.20	0.403	109
Major: Other	0.27	0.444	109
Fall term 2021	0.44	0.499	109
Spring term 2022	0.56	0.499	109
Response rates	Total	Share	Obs.
Rated both mentors	78	0.72	109
Rated any mentor	109	1.00	109

NOTE.— 'Overall' calculated as the average of Helpful, Fast, and Knowledgeable (pairwise deletion). See Appendix 6 (Table A1) for summary statistics for female and male students separately.

[11] One might worry that the high ratings circumscribe the study's external validity. However, gender bias at the top of the teacher rating distribution is arguably a relevant concern for young teachers aspiring to become tenured professors.

5.2 Preregistered analysis

In keeping with the preregistered analysis plan, our preferred model is the control-adjusted model without fixed effects (i.e., `covs`). Table 3 presents the study's main results. The models in Table 3 control for student gender, educational program, and age group and show that the teacher ratings are not biased against the female mentor. The lack of female bias does not reflect opposing effects across genders or different quality dimensions. The effects are insignificant and close to zero for both female and male students and similar across all dimensions.

Table 3. *Main effects*

	Overall	Helpful	Knowledge	Fast
	All students			
Female mentor	-0.0184	0.0188	-0.0359	-0.0408
	(0.0478)	(0.0588)	(0.0459)	(0.0622)
Observations	187	186	182	179
Student obs.	109	108	106	105
Mean rating	4.621	4.645	4.687	4.514
Std. dev. rating	0.689	0.737	0.702	0.810
	Female students			
Female mentor	0.00318	0.0277	-0.0350	0.0168
	(0.0516)	(0.0621)	(0.0409)	(0.0752)
Observations	123	122	119	117
Student obs.	72	71	69	69
Mean rating	4.738	4.770	4.782	4.641
Std. dev. rating	0.486	0.494	0.472	0.636
	Male students			
Female mentor	-0.0404	-0.00297	-0.0192	-0.112
	(0.0958)	(0.125)	(0.101)	(0.0973)
Observations	64	64	63	62
Student obs.	37	37	37	36
Mean rating	4.396	4.406	4.508	4.274
Std. dev. rating	0.486	0.494	0.472	0.636

NOTE.— Dependent variables defined column header. All cells represent separate regressions. All regressions include controls for student gender, educational program, and age. "Overall" is the average rating across the three dimensions 'Helpful', 'Knowledgeable', and 'Fast' (possible values 1-5). See Table A5 in Appendix 6 for the exact formulation of the SET questions.

The appendix contains results using unadjusted t-tests (Table A2) as well as fixed-effects models (Table A3). We also consider a different imputation method for the 'Overall' measure (not preregistered); see Table A4 in Appendix 6. These additional models confirm the baseline (zero) results. As

expected, the models with controls (i.e., the preferred model) and fixed effects are the most precise.

The concentration of evaluation scores around 4-5 (see Figure A2) implies that the standard deviation of teacher ratings is lower than we expected in the power analysis, meaning that the confidence bounds of the estimated treatment effect are tighter than expected. The point estimate on the average rating ('Overall') is -0.0184, which is 0.4 percent relative to the mean rating or 2.7 percent of a standard deviation. At the 95-percent level, the lower end of the confidence interval is -0.112, about 2.4 percent of the average rating or 15.9 percent of a standard deviation.

In conclusion, our analysis's failure to reject the null hypothesis of no bias does not reflect imprecise estimates relative to the existing studies. Figure 1 normalizes the treatment effect of this study and compares it with the most relevant previous studies (as well as across the three methods employed and across the three dimensions considered). Put in perspective, our 95%-confidence bounds reject the point estimates of other papers. Specifically, we reject effect sizes that are about 20 percent of the overall effect size in MacNell et al. (2015).

Figure 1. Female-Male Response Difference: Effect size comparison with previous literature

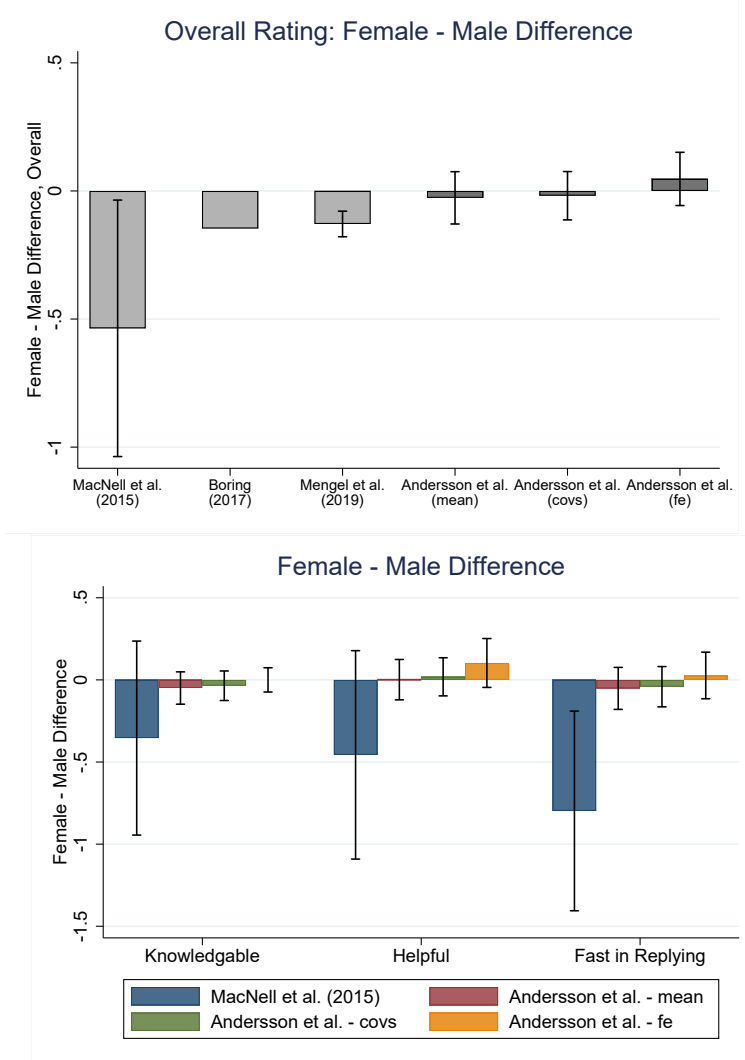

NOTE.— Female-Male Response Difference: Effect size comparison with previous literature. The overall rating for MacNell et al. (2015) is based on the three dimensions considered in this paper, computed by us using the original data in MacNell et al. (2015). Boring (2017) report effect sizes for female and male students and across the spring and fall terms separately. For comparison, the estimate for Boring (2017) reproduced here is the average difference from Table 2 in Boring (2017), weighted by sample size and rescaled from a 4-point scale to a 5-point scale. The bar and confidence band for Mengel et al. (2019) is based on the "Instructor-related"-rating in Table 5 in Mengel et al. (2019), without the interaction term for female student/teacher (generously provided by the authors). Andersson et al. refer to our estimates in this paper; `mean` is the unadjusted mean comparison, `covs` is our preferred model with controls, and `fe` is the fixed-effects model.

5.3 Explorative analysis of follow-up questions (not preregistered)

A complementary research question is if students expect female mentors to provide more help than male mentors. Such expectations risk underestimating the rating bias against females in observational studies if it means female teachers must work harder to earn the same rating. A possibility in our context is that the female mentor received more follow-up questions at a given response quality (see Appendix A.3 for an example). The main research question (SET ratings, analyzed in the previous section) is robust to such dynamics because all student inquiries, including follow-up questions, were randomly assigned to the female or male mentor. In less structured settings, however, a specific mentor would more likely answer their own follow-up questions.

Our data on email interactions can be used to investigate these mechanisms. We define follow-up questions as replies to a mentor's response within each subject. Using this definition, 275 mentor responses got a follow-up email from the students. The follow-up emails were usually new questions, but students sometimes replied to confirm that they understood the mentor's explanation. Sometimes, the students simply thanked the mentor for their help.

We define treatment as the gender signature of the previous mentor email and our first test is if the female mentor triggered more follow-up questions. Columns 1-3 of Table 4 show no difference between the male and female mentors in this regard. The female mentor did not receive more follow-up questions from students. We next consider the politeness of the follow-up emails, conditional upon a follow-up email having been sent. We used ChatGPT to determine the politeness of the student's email on a scale from 1 to 10 (see Appendix A.2 for the exact prompt used). Overall, the students were very polite – there were no negative emails; the chatbot's average politeness score is approximately 7.5. We validated the chatbot's politeness score using a more straightforward measure of politeness – an indicator variable equal to 1 if the student replied with a "thank you" (the correlation between the politeness score and the "thank you"-indicator is $0.3, t = 19.91$). Columns 4-9 in Table 4 show that the mentor's perceived gender did not affect the politeness of the students' follow-up questions.

Table 4. *Analysis of students' email interactions with the female teachers*

	Follow-up question			Politeness index			"Thank you"		
	All students	Female students	Male students	All students	Female students	Male students	All students	Female students	Male students
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Female mentor	0.00355	0.00461	0.000989	0.0105	-0.0806	0.126	0.0905	0.0596	0.135
	(0.0329)	(0.0421)	(0.0538)	(0.139)	(0.172)	(0.233)	(0.0576)	(0.0819)	(0.0755)
Constant	0.358	0.353	0.365	7.555	7.542	7.574	0.511	0.530	0.481
	(0.0288)	(0.0341)	(0.0516)	(0.123)	(0.127)	(0.246)	(0.0508)	(0.0674)	(0.0769)
Observations	765	453	312	275	161	114	275	161	114

NOTE.— Standard errors in parenthesis, clustered at the email address-level (203 clusters in total). The treatment variable 'Female mentor' is equal to 1 if the previous mentor was female and 0 otherwise; the 'Constant' refers to the intercept of the estimation. In models (1)-(3), the outcome variable is equal to 1 if the student sent a follow-up email to the mentor, zero otherwise. A 'follow-up email' is a new email with the same subject line from the same email address. In models (4)-(6), the outcome is an index between 1-10 measuring how polite the follow-up email was (using ChatGPT to determine the value). In models (7)-(9), the outcome is a dummy equal to 1 if the word "Tack" occurs in the first 100 characters of the student's follow-up email, zero otherwise. Models (4)-(9) are restricted to mentor emails that got a follow-up email. The correlation coefficient between the politeness index and the "thank you" indicator is 0.3050.

6 Conclusions

This study randomly assigned female and male names to online teacher correspondence and compared how students interacted with and rated these teachers in course evaluations. With this setup, the actual performance of the male and female teachers was balanced on average, meaning different ratings can be interpreted as student biases.

We find no evidence that course evaluations are biased against female teachers. Moreover, using our summary measure, the 95-percent confidence interval rejects effect sizes found in the previous literature. Design-wise, we also want to highlight that, in contrast to previous literature, our experiment was a double-blind, randomized natural field experiment. We confirm that the quality of the emails sent by the instructors was balanced across the female and male teachers and find no evidence that female teachers face a more demanding teaching environment in the setting studied here.

A remaining question is whether the external validity of our study extends to other settings or different subject pools. The educational environment studied is certainly typical of tertiary education. However, we note that our sample of students may differ from students in other disciplines or countries. Scandinavian countries are widely considered to have higher overall gender equality, which can potentially explain why our experiment finds no evidence of SET bias against women.

Ending on a more speculative note, one obvious avenue for future research is to increase the gender dosage over and above the sender name in text-based teaching. However, it is hard to imagine such an experimental setup that does not also introduce the issues inherent in audit studies, i.e. experimenter demand effects and less tight control of all relevant aspects other than gender. Recent technological advancements provide a potential, albeit speculative, solution to the problem. In the near future, one could imagine artificial teacher avatars, indistinguishable from real humans, in a remote teaching environment. In such a scenario, the dosage of gender (and other potential sources of bias) could be increased substantially while still keeping tight control over objective teacher quality.

References

Adams, Sophie, Sheree Bekker, Yanan Fan, Tess Gordon, Laura J Shepherd, Eve Slavich, and David Waters (2021). "Gender bias in student evaluations of teaching:'Punish [ing] those who fail to do their gender right'". *Higher Education*, 1–21.

Ahmed, Ali, Mark Granberg, and Shantanu Khanna (2021). "Gender discrimination in hiring: An experimental reexamination of the Swedish case". *PloS one* 16.1, e0245513.

Allport, Gordon Willard (1954). *The nature of prejudice*. Addison-wesley Reading, MA.

Aragón, Oriana R, Evava S Pietri, and Brian A Powell (2023). "Gender bias in teaching evaluations: The causal role of department gender composition". *Proceedings of the National Academy of Sciences* 120.4, e2118466120.

Arbuckle, Julianne and Benne D Williams (2003). "Students' perceptions of expressiveness: Age and gender effects on teacher evaluations". *Sex Roles* 49, 507–516.

Arceo-Gomez, Eva O and Raymundo M Campos-Vazquez (2019). "Gender stereotypes: The case of MisProfesores. com in Mexico". *Economics of Education Review* 72, 55–65.

Avdelningen för kvalitetsutveckling (Mar. 2020). "Vägledning för kursvärderingar: Stöd och råd i kursvärderingsarbetet". Tech. rep. UFV 2020/2340. Uppsala University.

Ayllón, Sara (2022). "Online teaching and gender bias". *Economics of Education Review* 89, 102280.

Bates, Laura (2015). "Female academics face huge sexist bias–no wonder there are so few of them". *The Guardian* 14.

Benton, Stephen L and William E Cashin (2013). "Student ratings of instruction in college and university courses". *Higher education: Handbook of theory and research: Volume 29*. Springer, 279–326.

Bertrand, Marianne and Esther Duflo (2017). "Field experiments on discrimination". *Handbook of economic field experiments* 1, 309–393.

Bertrand, Marianne and Sendhil Mullainathan (2004). "Are Emily and Greg more employable than Lakisha and Jamal? A field experiment on labor market discrimination". *American economic review* 94.4, 991–1013.

Bettinger, Eric P and Bridget Terry Long (2005). "Do faculty serve as role models? The impact of instructor gender on female students". *American Economic Review* 95.2, 152–157.

Binderkrantz, Anne Skorkjær and Mette Bisgaard (2023). "A gender affinity effect: The role of gender in teaching evaluations at a Danish university". *Higher Education*, 1–20.

Binderkrantz, Anne Skorkjær, Mette Bisgaard, and Berit Lassesen (2022). "Contradicting findings of gender bias in teaching evaluations: Evidence from two experiments in Denmark". *Assessment & Evaluation in Higher Education* 47.8, 1345–1357.

Boehmer, Devin M and William C Wood (2017). "Student vs. faculty perspectives on quality instruction: Gender bias, "hotness", and "easiness" in evaluating teaching". *Journal of Education for Business* 92.4, 173–178.

Boring, Anne (2017). "Gender biases in student evaluations of teaching". *Journal of public economics* 145, 27–41.

Boring, Anne and Kellie Ottoboni (2016). "Student evaluations of teaching (mostly) do not measure teaching effectiveness". *ScienceOpen Research*.

Buckles, Kasey (2019). "Fixing the leaky pipeline: Strategies for making economics work for women at every stage". *Journal of economic perspectives* 33.1, 43–60.

Carlsson, Magnus (2011). "Does hiring discrimination cause gender segregation in the Swedish labor market?" *Feminist Economics* 17.3, 71–102.

Carlsson, Magnus and Stefan Eriksson (2019). "In-group gender bias in hiring: Real-world evidence". *Economics Letters* 185, 108686.

Carrell, Scott E, Marianne E Page, and James E West (2010). "Sex and science: How professor gender perpetuates the gender gap". *The Quarterly journal of economics* 125.3, 1101–1144.

Ceci, Stephen J, Shulamit Kahn, and Wendy M Williams (2023). "Exploring gender bias in six key domains of academic science: An adversarial collaboration". *Psychological Science in the Public Interest* 24.1, 15–73.

Cheryan, Sapna, Sianna A Ziegler, Amanda K Montoya, and Lily Jiang (2017). "Why are some STEM fields more gender balanced than others?" *Psychological bulletin* 143.1, 1.

Fan, Yanan, Laura J Shepherd, Eve Slavich, David Waters, M Stone, Rachel Abel, and Emma L Johnston (2019). "Gender and cultural bias in student evaluations: Why representation matters". *PloS one* 14.2, e0209749.

Fox, Jesse and Wai Yen Tang (2017). "Women's experiences with general and sexual harassment in online video games: Rumination, organizational responsiveness, withdrawal, and coping strategies". *New media & society* 19.8, 1290–1307.

Hamermesh, Daniel S and Amy Parker (2005). "Beauty in the classroom: Instructors' pulchritude and putative pedagogical productivity". *Economics of Education Review* 24.4, 369–376.

Heffernan, Troy (2021). "Sexism, racism, prejudice, and bias: A literature review and synthesis of research surrounding student evaluations of courses

and teaching". *Assessment & Evaluation in Higher Education* 47.1, 144–154.

Hoffmann, Florian and Philip Oreopoulos (2009). "A professor like me the influence of instructor gender on college achievement". *Journal of human resources* 44.2, 479–494.

Jacob, Brian, Brian McCall, and Kevin Stange (2018). "College as country club: Do colleges cater to students' preferences for consumption?" *Journal of Labor Economics* 36.2, 309–348.

Kahn, Shulamit (1993). "Gender differences in academic career paths of economists". *The American Economic Review* 83.2, 52–56.

Kahn, Shulamit (1995). "Women in the economics profession". *Journal of Economic Perspectives* 9.4, 193–205.

Karlsson, Martin and Erik Lundberg (2012). "I betraktarens ögon–Betydelsen av kön och ålder för studenters läraromdömen". *Högre utbildning* 2.1, 19–32.

Keng, Shao-Hsun (2018). "Tenure system and its impact on grading leniency, teaching effectiveness and student effort". *Empirical Economics* 55, 1207–1227.

Khazan, E, J Borden, S Johnson, and L Greenhaw (2019). "Examining Gender Bias in Student Evaluations of Teaching for Graduate Teaching Assistants". *NACTA Journal* 64, 422–427.

Kreisman, Daniel and Jonathan Smith (2023). "Distinctively Black names and educational outcomes". *Journal of Political Economy* 131.4, 877–897.

Kulik, James A (2001). "Student ratings: Validity, utility, and controversy". *New directions for institutional research* 2001.109, 9–25.

Lundberg, Shelly and Jenna Stearns (2019). "Women in economics: Stalled progress". *Journal of Economic Perspectives* 33.1, 3–22.

MacNell, Lillian, Adam Driscoll, and Andrea N Hunt (2015). "What's in a name: Exposing gender bias in student ratings of teaching". *Innovative Higher Education* 40.4, 291–303.

Mengel, Friederike, Jan Sauermann, and Ulf Zölitz (2019). "Gender bias in teaching evaluations". *Journal of the European economic association* 17.2, 535–566.

Moore, Melanie and Richard Trahan (1998). "Tenure status and grading practices". *Sociological Perspectives* 41.4, 775–781.

Moss-Racusin, Corinne A, John F Dovidio, Victoria L Brescoll, Mark J Graham, and Jo Handelsman (2012). "Science faculty's subtle gender biases favor male students". *Proceedings of the national academy of sciences* 109.41, 16474–16479.

Neumark, David, Roy J Bank, and Kyle D Van Nort (1996). "Sex discrimination in restaurant hiring: An audit study". *The Quarterly journal of economics* 111.3, 915–941.

Nielsen, Mathias W (2016). "Limits to meritocracy? Gender in academic recruitment and promotion processes". *Science and Public Policy* 43.3, 386–399.

Özgümüs, Asri, Holger A Rau, Stefan T Trautmann, and Christian König-Kersting (2020). "Gender bias in the evaluation of teaching materials". *Frontiers in psychology* 11, 1074.

Paredes, Valentina (2014). "A teacher like me or a student like me? Role model versus teacher bias effect". *Economics of Education Review* 39, 38–49.

Paredes, Valentina, M. Daniele Paserman, and Francisco J. Pino (July 2023). "Does Economics Make You Sexist?" *The Review of Economics and Statistics*, 1–47. eprint: https://direct.mit.edu/rest/article-pdf/doi/10.1162/rest_a_01356/2150372/rest_a_01356.pdf.

Petit, Pascale (2007). "The effects of age and family constraints on gender hiring discrimination: A field experiment in the French financial sector". *Labour Economics* 14.3, 371–391.

Reid, Landon D (2010). "The role of perceived race and gender in the evaluation of college teaching on RateMyProfessors. com." *Journal of Diversity in higher Education* 3.3, 137.

Remmers, Herman H and George C Brandenburg (1927). "Experimental data on the Purdue rating scale for instructors". *Educational Administration and Supervision* 13.6, 399–406.

Riach, Peter A and Judith Rich (2006). "An experimental investigation of sexual discrimination in hiring in the English labor market". *The BE Journal of Economic Analysis & Policy* 6.2.

Rosen, Andrew S (2018). "Correlations, trends and potential biases among publicly accessible web-based student evaluations of teaching: a large-scale study of RateMyProfessors.com data". *Assessment & Evaluation in Higher Education* 43.1, 31–44.

Rothstein, Donna S (1995). "Do female faculty influence female students' educational and labor market attainments?" *ILR Review* 48.3, 515–530.

Stroebe, Wolfgang (2020). "Student evaluations of teaching encourages poor teaching and contributes to grade inflation: A theoretical and empirical analysis". *Basic and Applied Social Psychology* 42.4, 276–294.

Stuber, Jenny M, Amanda Watson, Adam Carle, and Kristin Staggs (2009). "Gender expectations and on-line evaluations of teaching: Evidence from RateMyProfessors. com". *Teaching in Higher Education* 14.4, 387–399.

Tangalakis, Kathy, Kate Kelly, Natalie KonYu, and Dianne Hall (2022). "The impact of teaching from home during the covid-19 pandemic on the student evaluations of female academics". *Journal of University Teaching & Learning Practice* 19.1, 160–175.

Uttl, Bob and Victoria C Violo (2021). "Small samples, unreasonable generalizations, and outliers: Gender bias in student evaluation of teaching or three unhappy students?" *ScienceOpen Research*.

Weisshaar, Katherine (2017). "Publish and perish? An assessment of gender gaps in promotion to tenure in academia". *Social Forces* 96.2, 529–560.

Wong, Jennifer S and Jessica Bouchard (2021). "Are students gender-neutral in their assessment of online teaching staff?" *Assessment & Evaluation in Higher Education* 46.5, 719–739.

Wu, Alice H (2018). "Gendered language on the economics job market rumors forum". *AEA Papers and Proceedings*. Vol. 108. American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203, 175–179.

Appendix A: Additional figures and tables

Figure A1. Timeline of Emails

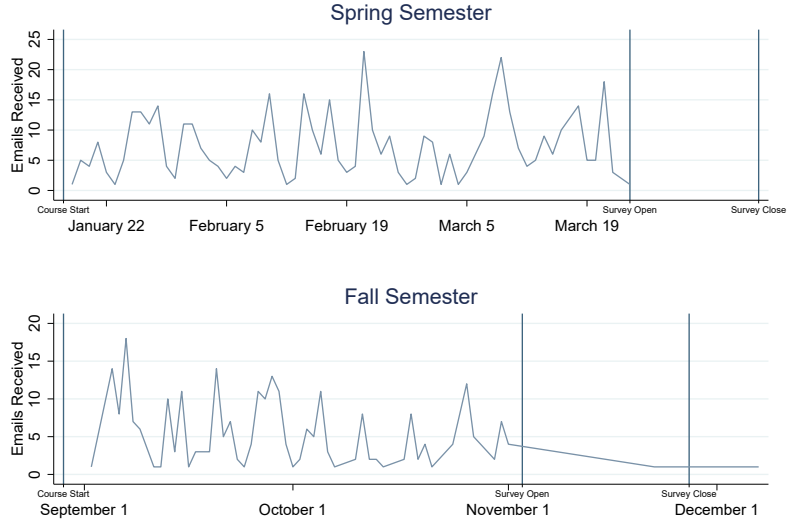

NOTE.— Number of emails received during the spring and fall semester, respectively. Start of course, opening and closing of the course evaluation indicated by vertical lines.

Figure A2. Survey response distribution

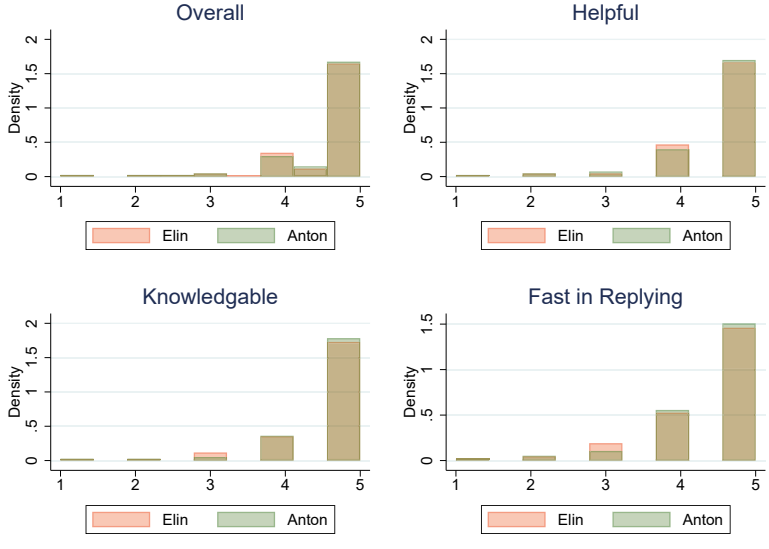

NOTE.— Male and Female teacher. Response distribution for male and female instructor for all dimension and the unweighted mean 'Overall'.

Table A1. *Summary statistics: Course evaluation*

	Female students			Male students		
	Mean	Std. dev	Obs.	Mean	Std. dev	Obs.
Ratings (Male mentor)						
Helpful	4.76	0.567	59	4.42	0.992	31
Fast in replying	4.64	0.645	56	4.34	1.010	29
Knowledgeable	4.81	0.441	57	4.53	0.973	30
Overall (average)	4.74	0.509	59	4.43	0.895	31
Ratings (Female mentor)						
Helpful	4.78	0.419	63	4.39	1.059	33
Fast in replying	4.64	0.633	61	4.21	1.053	33
Knowledgeable	4.76	0.502	62	4.48	1.004	33
Overall (average)	4.73	0.468	64	4.36	0.973	33
Student characteristics						
<20 years old	0.03	0.165	72	0.05	0.229	37
20 years old	0.21	0.409	72	0.22	0.417	37
21 years old	0.26	0.444	72	0.22	0.417	37
22 years old	0.18	0.387	72	0.16	0.374	37
23 years old	0.10	0.298	72	0.11	0.315	37
24 years old	0.13	0.333	72	0.11	0.315	37
>24 years	0.10	0.298	72	0.14	0.347	37
Major: Business	0.50	0.504	72	0.59	0.498	37
Major: Pol. Sci.	0.26	0.444	72	0.08	0.277	37
Major: Other	0.24	0.428	72	0.32	0.475	37
Fall term 2021	0.33	0.475	72	0.65	0.484	37
Spring term 2022	0.67	0.475	72	0.35	0.484	37
Response rates	Total	Share	Obs.	Total	Share	Obs.
Rated both mentors	51	0.71	72	27	0.73	37
Rated any mentor	72	1.00	72	37	1.00	37

NOTE.— This table is the same as the Summary Statistics table (Table 2), but over male and female students. 'Overall' is calculated as the average of Helpful, Fast, and Knowledgeable (pairwise deletion).

Table A2. *Unadjusted treatment effects (no control variables)*

	Overall	Helpful	Knowledge	Fast
	All students			
Female mentor	-0.0269	0.00139	-0.0495	-0.0518
	(0.0521)	(0.0627)	(0.0502)	(0.0652)
Observations	187	186	182	179
Student obs.	109	108	106	105
Mean rating	4.621	4.645	4.687	4.514
Std. dev. rating	0.689	0.737	0.702	0.810
	Female students			
Female mentor	-0.00856	0.0151	-0.0490	-0.00351
	(0.0540)	(0.0648)	(0.0446)	(0.0780)
Observations	123	122	119	117
Student obs.	72	71	69	69
Mean rating	4.738	4.770	4.782	4.641
Std. dev. rating	0.486	0.494	0.472	0.636
	Male students			
Female mentor	-0.0665	-0.0254	-0.0485	-0.133
	(0.109)	(0.132)	(0.115)	(0.115)
Observations	64	64	63	62
Student obs.	37	37	37	36
Mean rating	4.396	4.406	4.508	4.274
Std. dev. rating	0.929	1.019	0.982	1.027

NOTE.— Dependent variables defined column header. All cells represent separate regressions. Unlike the main table (Table 3) this table presents regression estimates without control variables. 'Overall' is the average rating across the three dimensions 'Helpful,' 'Knowledgeable,' and 'Fast' (possible values 1-5). See A5 for the exact formulation of the SET questions.

Table A3. *Fixed effects analysis*

	Overall	Helpful	Knowledge	Fast
	All students			
Female mentor	0.0470	0.103	0	0.0270
	(0.0530)	(0.0759)	(0.0377)	(0.0723)
Observations	156	156	153	149
Student obs.	78	78	77	75
Mean grade	4.630	4.654	4.699	4.523
	Female students			
Female mentor	0.0523	0.0980	0	0.0625
	(0.0694)	(0.0912)	(0.0406)	(0.0984)
Observations	102	102	100	96
Student obs.	51	51	50	48
Mean grade	4.745	4.775	4.780	4.656
	Male students			
Female mentor	0.0370	0.111	0	-0.0385
	(0.0817)	(0.139)	(0.0799)	(0.0973)
Observations	54	54	53	53
Student obs.	27	27	27	27
Mean grade	4.414	4.426	4.547	4.283

NOTE.— Dependent variables defined column header. All cells represent separate regressions. Unlike the main table (Table 3) this table presents regression estimates with fixed student effects. The sample size is lower in this table because we restrict the estimation to students who rate both Elin and Anton. Overall is the average rating across the three dimensions helpful, knowledge, and fast (possible values 1-5). See A5 for the exact formulation of the SET questions. Treatment effect is exactly zero for the outcome Knowledge because no student rated Elin and Anton differently in this dimension.

Table A4. *Alternative outcome: imputed Overall.*

	All students	Male students	Female students
	Unadjusted means		
Female mentor	-0.0321	-0.0742	-0.0123
	(0.0514)	(0.106)	(0.0538)
Observations	187	64	123
Student obs.	109	37	72
Mean grade	4.615	4.556	4.666
	With student controls		
Female mentor	-0.0241	-0.0481	0.000679
	(0.0473)	(0.0928)	(0.0512)
Observations	187	64	123
Student obs.	109	37	72
Mean grade	4.615	4.556	4.666
	Fixed effects estimation		
Female mentor	0.0419	0.0222	0.0523
	(0.0568)	(0.0815)	(0.0761)
Observations	187	64	123
Student obs.	109	37	72
Mean grade	4.625	4.421	4.734

NOTE.— All cells represent separate regressions. Student Controls include gender, educational program, and age. Overall is the average rating across the three dimensions helpful, knowledge, and fast (possible values 1-5).

Table A5. *SET Questions*

11.	Was mail mentor Elin helpful?				
Not at all				To a very large extent	Don't know
○	○	○	○	○	○
12.	Was mail mentor Anton helpful?				
Not at all				To a very large extent	Don't know
○	○	○	○	○	○
13.	Was mail mentor Elin knowledgable?				
Not at all				To a very large extent	Don't know
○	○	○	○	○	○
14.	Was mail mentor Anton knowledgable?				
Not at all				To a very large extent	Don't know
○	○	○	○	○	○
15.	Did mail mentor Elin answer emails quickly?				
Not at all				To a very large extent	Don't know
○	○	○	○	○	○
16.	Did mail mentor Anton answer emails quickly?				
Not at all				To a very large extent	Don't know
○	○	○	○	○	○

NOTE.— Table displays the questions used to evaluate the male (Anton) and female (Elin) instructors (translated from Swedish). Students were randomly allocated surveys in the displayed order or the reverse, where questions about the male instructor (Anton) preceded questions about the female instructor (Elin).

A.1 Power calculation

The required sample size n (one sample, continuous outcome) is given by $n = ((Z_(1-\alpha/2) + Z_(1-\beta))/\text{ES})^2$. With $\alpha = 0.05$ (test at 5 percent significance level) and $\beta = 0.2$ (power at 80 percent level) we get $Z_(1-\alpha/2) = Z_{0.975} = 1.96$ and $Z_{(1-\beta)} = Z_{0.8} = 0.84$. The ES $\equiv |\mu|/\sigma$, where ES stands for effect size, μ is the mean difference between "male" and "female" average scores in the SET. σ is the standard deviation of the test statistic. Based on previous literature we set $\mu = 0.2$ (on a 5 grade scale). In order to estimate σ we make use of previous SETs (on the same course) and chose two questions that have an average difference of $\mu \approx 0.2$. This difference then has an estimated standard deviation of $\hat{\sigma} = 0.64$. Based on these assumptions the required sample size is:

$$n = \left(\frac{1.96 + 0.84}{0.2/0.64}\right)^2 = 80.282 \tag{A1}$$

A.2 Instructions to the ChatBot

We used Chat GPT to rate the tone of the students' responses. Chat GPT is an artificial intelligence chatbot developed by OpenAI.

The prompt was as follows. In Swedish (original):

Kan du bedöma dessa mail utifrån hur trevlig ton det är i dem? Det är XX st, numrerade i stigande ordning (alla mail har nummer över 1000000 så du kan anta att ett nytt mail börjat och det förra avslutats när det kommer ett nummer över 1000000).
Kan du för varje mail göra en numerisk bedömning på en skala 1-10 där 1 är mycket otrevligt och 10 är extremt trevligt (Svara på formen 1000011:8 etc utan ytterligare kommentarer och med radbrytning mellan varje svar). Här kommer mailen:

Translation:

Can you rate these emails based on how pleasant the tone is? There are XX emails, numbered in ascending order (all emails have numbers over 1000000 so you can assume that a new email has started and the previous one has ended when a number over 1000000 arrives).
For each email, can you make a numerical assessment on a scale of 1-10 where 1 is very unpleasant and 10 is extremely nice (Reply in the form 1000011:8 etc without further comments and with line breaks between each answer). Here comes the emails:

A.3 Correspondence examples

In this appendix section, we use a real email conversation to illustrate how the correspondence between the students and the mentors could look. We first present the original in Swedish and then an English translation. The emails are presented here just like they appear in the students' email client (starting with the last email and then the email subject's history appended). Note that in the Swedish original, the meta-data of the email is sometimes written in English. This is not a typo. When pressing "Reply", the student's email client creates the meta-data, and if the student's client language is set to English, the date and the field names of the mail mentor's first response will be in the English format (i.e., "From:", "To:" etc., instead of "Från:", "Till:").

The emails are unedited, except that integrity-compromising details (like the student's name) have been changed.

Correspondence example, in Swedish

Från: Anna Andersson <Anna.Andersson.xxxx@student.uu.se>
Skickat: den 31 januari 2022 13:56
Till: Mailmentor <mailmentor@nek.uu.se>
Ämne: Sv: Metod för partiell derivata

Hej,
det var rätt tolkat! Tack ska du ha, Anton!
//Anna

Från: Mailmentor <mailmentor@nek.uu.se>
Skickat: den 31 januari 2022 08:50
Till: Anna Andersson <Anna.Andersson.xxxx@student.uu.se>
Ämne: Sv: Metod för partiell derivata

Hej Anna!

Jag är inte helt säker på vad du frågar efter men jag ger det ett försök så kan du rätta mig i efterhand. När man tar partialderivatan så är det viktigt att tydligt skriva vilken variabel som man deriverar med avseende på. I fråga 4 på räkneövningen så har vi efterfrågefunktionen Q och i den finns två priser, priset för den huvudsakliga varan (den vars Q vi räknar) och en annan vara på marknaden. Det finns ingen direkt "lag" på hur det ska se ut men jag skulle skriva Q_p (alltså p i index/subscript) för priselasticitet och Q_{pj} (alltså dubbelindex i princip) för korspriselasticitet.

Som sagt, om detta inte var det du frågade efter så hör av dig igen bara!

Med vänliga hälsningar,
Anton

Från: Anna Andersson <Anna.Andersson.xxxx@student.uu.se>
Skickat: den 28 januari 2022 11:54
Till: Mailmentor <mailmentor@nek.uu.se>
Ämne: Sv: Metod för partiell derivata

Hej Elin,

tack så mycket för en ingående förklaring!

När man använder sig av partialderivata för att räkna ut priselasticitet, hur använder jag mig av tecknet för partialderivata? Vad "översätts" den till, om du förstår vad jag menar?

Min referens är från fråga 4 från dagens matematikövning i NEK A.
//Anna

Hämta Outlook för Android

From: Mailmentor <mailmentor@nek.uu.se>
Sent: Friday, January 28, 2022 10:59:40 AM
To: Anna Andersson <Anna.Andersson.xxxx@student.uu.se>
Subject: Sv: Metod för partiell derivata

Hej Anna!

Absolut, jag ger det ett försök.

Partialderivator använder vi i de fall när våra funktioner innehåller mer än en variabel, exempelvis $f(x,y)$. Partialderivatan hjälper oss att se hur funktionens värde, exempelvis en nyttofunktion, ändras om vi bara ändrar värdet på en av våra variabler och låter den andra vara låst vid ett värde (ett värde som vi inte känner till men vi antar att den är fix). Exempelvis kanske vi är intresserade av att se hur en persons nyttofunktion ändras om man skiftar hur mycket den konsumerar av produkt A, givet att vi inte kan ändra på hur mycket av produkt B den kan/får konsumera.

Ett företags vinstfunktion, som beror dels av hur många produkter de säljer och dels av vilket pris de sätter, Vinst $= f(p,q)$. Då kan det vara intressant att se hur mycket vinsten ändras av att ändra priset men låta kvantiteten vara fast (det kanske är svårt eller tar tid att lägga om produktionen).

När man tar partialderivatan av en funktion $f(x,y)$ så gäller alla samma regler som när man tar en vanlig derivata utom det faktum att den variabel som man inte deriverar med avseende på behandlas som en konstant.

Se fråga 81a bifogat. Det här är ett enkelt exempel hur man räknar ut partialderivatan. När vi deriverar $f(x,y)$ med avseende på x så hänger y bara med som en konstant, precis som i vanliga fall när man har något multiplicerat med den variabel som deriverar på, och vi deriverar x^2 som vanligt.

När vi deriverar med avseende på y så blir endast x^2 kvar, precis på samma sätt som bara 5 hade varit kvar om vi deriverade funktion $f(y) = 5y$. Det spelar ingen roll att x är upphöjt till något för den behandlas som en helt vanlig konstant.

81.a)

$$\begin{aligned} f(x,y) &= x^2y \\ f'_x &= 2xy \\ f'_y &= x^2 \end{aligned}$$

Hoppas det här gjorde det något klarare!

Med vänliga hälsningar,
Elin

Från: Anna Andersson <Anna.Andersson.xxxx@student.uu.se>
Skickat: den 28 januari 2022 10:49
Till: Mailmentor <mailmentor@nek.uu.se>
Ämne: Metod för partiell derivata

Hej Mailmentor!

Jag har varit frånvarande och behöver ta igen material som jag har missat.

Jag har en fråga angående partiell derivata. Jag har problem med att förstå hur den fungerar och vad det är som räknas ut med hjälp av den.

Skulle jag kunna få ett räkneexempel, steg för steg, på partiell derivatas användningsområde och metod?

Tackar på förhand,

Correspondence, English translation

From: Anna Andersson <Anna.Andersson.xxxx@student.uu.se>
Sent: January 31, 2022 1:56 PM
To: Mailmentor <mailmentor@nek.uu.se>
Subject: Re: Method for partial derivative

Hi,

that was correctly interpreted! Thank you, Anton!

//Anna

From: Mailmentor <mailmentor@nek.uu.se>
Sent: January 31, 2022 8:50 AM
To: Anna Andersson <Anna.Andersson.xxxx@student.uu.se>
Subject: Re: Method for partial derivative

Hi Anna!

I'm not entirely sure what you're asking for, but I'll give it a try and you can correct me afterwards. When taking the partial derivative, it is important to clearly indicate which variable you are differentiating with respect to. In question 4 of the math exercise, we have the demand function Q, which contains two prices: the price of the main product (whose Q we are calculating) and the price of another product on the market. There is no direct "rule" for how this should be written, but I would write Q_p (thus p in subscript) for price elasticity and Q_{pj} (thus double subscript in principle) for cross-price elasticity.

If this wasn't what you were asking about, please let me know and ask again!

Best regards,
Anton

From: Anna Andersson <Anna.Andersson.xxxx@student.uu.se>
Sent: January 28, 2022 11:54 AM
To: Mailmentor <mailmentor@nek.uu.se>
Subject: Re: Method for partial derivative

Hi Elin,

thank you very much for a detailed explanation!

When using partial derivatives to calculate price elasticity, how do I use the sign for partial derivatives? What does it "translate" to, if you know what I mean?

My reference is from question 4 on today's math exercise in NEK A.
//Anna

Get Outlook for Android

From: Mailmentor <mailmentor@nek.uu.se>
Sent: January 28, 2022 10:59 AM
To: Anna Andersson <Anna.Andersson.xxxx@student.uu.se>
Subject: Re: Method for partial derivative

Hi Anna!

Certainly, I'll give it a try.

We use partial derivatives in cases where our functions contain more than one variable, such as $f(x,y)$. The partial derivative helps us see how the function value, such as a utility function, changes if we only change the value of one of our variables and leave the other fixed at a value (a value that we do not know but assume is fixed). For example, we might be interested in seeing how a person's utility function changes if we shift how much they consume of product A, given that we cannot change how much of product B they can consume.

A company's profit function, which depends partly on how many products they sell and partly on the price they set, Profit $= f(p,q)$. Then it can be interesting to see how much profit changes by changing the price but keeping the quantity fixed (it may be difficult or time-consuming to change the production).

When taking the partial derivative of a function $f(x,y)$, all the same rules apply as when taking a regular derivative, except for the fact that the variable that you are not differentiating with respect to is treated as a constant.

Please see question 81a attached. This is a simple example of how to calculate a partial derivative. When we differentiate $f(x,y)$ with respect to x, y simply hangs on as a constant, just like in regular cases when you have something multiplied by the variable you are differentiating with respect to, and we dif-

ferentiate x^2 as usual.

When we differentiate with respect to y, only x^2 remains, just as only 5 would have remained if we differentiated the function $f(y) = 5y$. It doesn't matter that x is raised to something since it is treated like a regular constant.

81.a)

$$\begin{aligned} f(x,y) &= x^2y \\ f'_x &= 2xy \\ f'_y &= x^2 \end{aligned}$$

Hope this made it clearer!

Best regards,
Elin

From: Anna Andersson <Anna.Andersson.xxxx@student.uu.se>
Sent: 28 January 2022 10:49
To: Mailmentor mailmentor@nek.uu.se
Subject: Method for Partial Derivatives

Hi Mailmentor!

I've been absent and need to catch up on material that I missed.

I have a question regarding partial derivatives. I have trouble understanding how it works and what it calculates.

Could I get a step-by-step calculation example on the usage and method of partial derivatives?

Thank you in advance,
Anna Andersson

Economic Studies

1987:1 Haraldson, Marty. To Care and To Cure. A linear programming approach to national health planning in developing countries. 98 pp.

1989:1 Chryssanthou, Nikos. The Portfolio Demand for the ECU. A Transaction Cost Approach. 42 pp.

1989:2 Hansson, Bengt. Construction of Swedish Capital Stocks, 1963-87. An Application of the Hulten-Wykoff Studies. 37 pp.

1989:3 Choe, Byung-Tae. Some Notes on Utility Functions Demand and Aggregation. 39 pp.

1989:4 Skedinger, Per. Studies of Wage and Employment Determination in the Swedish Wood Industry. 89 pp.

1990:1 Gustafson, Claes-Håkan. Inventory Investment in Manufacturing Firms. Theory and Evidence. 98 pp.

1990:2 Bantekas, Apostolos. The Demand for Male and Female Workers in Swedish Manufacturing. 56 pp.

1991:1 Lundholm, Michael. Compulsory Social Insurance. A Critical Review. 109 pp.

1992:1 Sundberg, Gun. The Demand for Health and Medical Care in Sweden. 58 pp.

1992:2 Gustavsson, Thomas. No Arbitrage Pricing and the Term Structure of Interest Rates. 47 pp.

1992:3 Elvander, Nils. Labour Market Relations in Sweden and Great Britain. A Comparative Study of Local Wage Formation in the Private Sector during the 1980s. 43 pp.

12 Dillén, Mats. Studies in Optimal Taxation, Stabilization, and Imperfect Competition. 1993. 143 pp.

13 Banks, Ferdinand E.. A Modern Introduction to International Money, Banking and Finance. 1993. 303 pp.

14 Mellander, Erik. Measuring Productivity and Inefficiency Without Quantitative Output Data. 1993. 140 pp.

15 Ackum Agell. Susanne. Essays on Work and Pay. 1993. 116 pp.

16 Eriksson, Claes. Essays on Growth and Distribution. 1994. 129 pp.

17 Banks, Ferdinand E.. A Modern Introduction to International Money, Banking and Finance. 2nd version, 1994. 313 pp.

18 Apel, Mikael. Essays on Taxation and Economic Behavior. 1994. 144 pp.

19 Dillén, Hans. Asset Prices in Open Monetary Economies. A Contingent Claims Approach. 1994. 100 pp.

20 Jansson, Per. Essays on Empirical Macroeconomics. 1994. 146 pp.

21 Banks, Ferdinand E.. A Modern Introduction to International Money, Banking, and Finance. 3rd version, 1995. 313 pp.

22 Dufwenberg, Martin. On Rationality and Belief Formation in Games. 1995. 93 pp.

23 Lindén, Johan. Job Search and Wage Bargaining. 1995. 127 pp.

24 Shahnazarian, Hovick. Three Essays on Corporate Taxation. 1996. 112 pp.

25 Svensson, Roger. Foreign Activities of Swedish Multinational Corporations. 1996. 166 pp.

26 Sundberg, Gun. Essays on Health Economics. 1996. 174 pp.

27 Sacklén, Hans. Essays on Empirical Models of Labor Supply. 1996. 168 pp.

28 Fredriksson, Peter. Education, Migration and Active Labor Market Policy. 1997. 106 pp.

29 Ekman, Erik. Household and Corporate Behaviour under Uncertainty. 1997. 160 pp.

30 Stoltz, Bo. Essays on Portfolio Behavior and Asset Pricing. 1997. 122 pp.

31 Dahlberg, Matz. Essays on Estimation Methods and Local Public Economics. 1997. 179 pp.

32 Kolm, Ann-Sofie. Taxation, Wage Formation, Unemployment and Welfare. 1997. 162 pp.

33 Boije, Robert. Capitalisation, Efficiency and the Demand for Local Public Services. 1997. 148 pp.

34 Hort, Katinka. On Price Formation and Quantity Adjustment in Swedish Housing Markets. 1997. 185 pp.

35 Lindström, Thomas. Studies in Empirical Macroeconomics. 1998. 113 pp.

36 Hemström, Maria. Salary Determination in Professional Labour Markets. 1998. 127 pp.

37 Forsling, Gunnar. Utilization of Tax Allowances and Corporate Borrowing. 1998. 96 pp.

38 Nydahl, Stefan. Essays on Stock Prices and Exchange Rates. 1998. 133 pp.

39 Bergström, Pål. Essays on Labour Economics and Econometrics. 1998. 163 pp.

40 Heiborn, Marie. Essays on Demographic Factors and Housing Markets. 1998. 138 pp.

41 Åsberg, Per. Four Essays in Housing Economics. 1998. 166 pp.

42 Hokkanen, Jyry. Interpreting Budget Deficits and Productivity Fluctuations. 1998. 146 pp.

43 Lunander, Anders. Bids and Values. 1999. 127 pp.

44 Eklöf, Matias. Studies in Empirical Microeconomics. 1999. 213 pp.

45 Johansson, Eva. Essays on Local Public Finance and Intergovernmental Grants. 1999. 156 pp.

46 Lundin, Douglas. Studies in Empirical Public Economics. 1999. 97 pp.

47 Hansen, Sten. Essays on Finance, Taxation and Corporate Investment. 1999. 140 pp.

48 Widmalm, Frida. Studies in Growth and Household Allocation. 2000. 100 pp.

49 Arslanogullari, Sebastian. Household Adjustment to Unemployment. 2000. 153 pp.

50 Lindberg, Sara. Studies in Credit Constraints and Economic Behavior. 2000. 135 pp.

51 Nordblom, Katarina. Essays on Fiscal Policy, Growth, and the Importance of Family Altruism. 2000. 105 pp.

52 Andersson, Björn. Growth, Saving, and Demography. 2000. 99 pp.

53 Åslund, Olof. Health, Immigration, and Settlement Policies. 2000. 224 pp.

54 Bali Swain, Ranjula. Demand, Segmentation and Rationing in the Rural Credit Markets of Puri. 2001. 160 pp.

55 Löfqvist, Richard. Tax Avoidance, Dividend Signaling and Shareholder Taxation in an Open Economy. 2001. 145 pp.

56 Vejsiu, Altin. Essays on Labor Market Dynamics. 2001. 209 pp.

57 Zetterström, Erik. Residential Mobility and Tenure Choice in the Swedish Housing Market. 2001. 125 pp.

58 Grahn, Sofia. Topics in Cooperative Game Theory. 2001. 106 pp.

59 Laséen, Stefan. Macroeconomic Fluctuations and Microeconomic Adjustments. Wages, Capital, and Labor Market Policy. 2001. 142 pp.

60 Arnek, Magnus. Empirical Essays on Procurement and Regulation. 2002. 155 pp.

61 Jordahl, Henrik. Essays on Voting Behavior, Labor Market Policy, and Taxation. 2002. 172 pp.

62 Lindhe, Tobias. Corporate Tax Integration and the Cost of Capital. 2002. 102 pp.

63 Hallberg, Daniel. Essays on Household Behavior and Time-Use. 2002. 170 pp.

64 Larsson, Laura. Evaluating Social Programs: Active Labor Market Policies and Social Insurance. 2002. 126 pp.

65 Bergvall, Anders. Essays on Exchange Rates and Macroeconomic Stability. 2002. 122 pp.

66 Nordström Skans, Oskar. Labour Market Effects of Working Time Reductions and Demographic Changes. 2002. 118 pp.

67 Jansson, Joakim. Empirical Studies in Corporate Finance, Taxation and Investment. 2002. 132 pp.

68 Carlsson, Mikael. Macroeconomic Fluctuations and Firm Dynamics: Technology, Production and Capital Formation. 2002. 149 pp.

69 Eriksson, Stefan. The Persistence of Unemployment: Does Competition between Employed and Unemployed Job Applicants Matter? 2002. 154 pp.

70 Huitfeldt, Henrik. Labour Market Behaviour in a Transition Economy: The Czech Experience. 2003. 110 pp.

71 Johnsson, Richard. Transport Tax Policy Simulations and Satellite Accounting within a CGE Framework. 2003. 84 pp.

72 Öberg, Ann. Essays on Capital Income Taxation in the Corporate and Housing Sectors. 2003. 183 pp.

73 Andersson, Fredrik. Causes and Labor Market Consequences of Producer Heterogeneity. 2003. 197 pp.

74 Engström, Per. Optimal Taxation in Search Equilibrium. 2003. 127 pp.

75 Lundin, Magnus. The Dynamic Behavior of Prices and Investment: Financial Constraints and Customer Markets. 2003. 125 pp.

76 Ekström, Erika. Essays on Inequality and Education. 2003. 166 pp.

77 Barot, Bharat. Empirical Studies in Consumption, House Prices and the Accuracy of European Growth and Inflation Forecasts. 2003. 137 pp.

78 Österholm, Pär. Time Series and Macroeconomics: Studies in Demography and Monetary Policy. 2004. 116 pp.

79 Bruér, Mattias. Empirical Studies in Demography and Macroeconomics. 2004. 113 pp.

80 Gustavsson, Magnus. Empirical Essays on Earnings Inequality. 2004. 154 pp.

81 Toll, Stefan. Studies in Mortgage Pricing and Finance Theory. 2004. 100 pp.

82 Hesselius, Patrik. Sickness Absence and Labour Market Outcomes. 2004. 109 pp.

83 Häkkinen, Iida. Essays on School Resources, Academic Achievement and Student Employment. 2004. 123 pp.

84 Armelius, Hanna. Distributional Side Effects of Tax Policies: An Analysis of Tax Avoidance and Congestion Tolls. 2004. 96 pp.

85 Ahlin, Åsa. Compulsory Schooling in a Decentralized Setting: Studies of the Swedish Case. 2004. 148 pp.

86 Heldt, Tobias. Sustainable Nature Tourism and the Nature of Tourists' Cooperative Behavior: Recreation Conflicts, Conditional Cooperation and the Public Good Problem. 2005. 148 pp.

87 Holmberg, Pär. Modelling Bidding Behaviour in Electricity Auctions: Supply Function Equilibria with Uncertain Demand and Capacity Constraints. 2005. 43 pp.

88 Welz, Peter. Quantitative new Keynesian macroeconomics and monetary policy 2005. 128 pp.

89 Ågren, Hanna. Essays on Political Representation, Electoral Accountability and Strategic Interactions. 2005. 147 pp.

90 Budh, Erika. Essays on environmental economics. 2005. 115 pp.

91 Chen, Jie. Empirical Essays on Housing Allowances, Housing Wealth and Aggregate Consumption. 2005. 192 pp.

92 Angelov, Nikolay. Essays on Unit-Root Testing and on Discrete-Response Modelling of Firm Mergers. 2006. 127 pp.

93 Savvidou, Eleni. Technology, Human Capital and Labor Demand. 2006. 151 pp.

94 Lindvall, Lars. Public Expenditures and Youth Crime. 2006. 112 pp.

95 Söderström, Martin. Evaluating Institutional Changes in Education and Wage Policy. 2006. 131 pp.

96 Lagerström, Jonas. Discrimination, Sickness Absence, and Labor Market Policy. 2006. 105 pp.

97 Johansson, Kerstin. Empirical essays on labor-force participation, matching, and trade. 2006. 168 pp.

98 Ågren, Martin. Essays on Prospect Theory and the Statistical Modeling of Financial Returns. 2006. 105 pp.

99 Nahum, Ruth-Aïda. Studies on the Determinants and Effects of Health, Inequality and Labour Supply: Micro and Macro Evidence. 2006. 153 pp.

100 Žamac, Jovan. Education, Pensions, and Demography. 2007. 105 pp.

101 Post, Erik. Macroeconomic Uncertainty and Exchange Rate Policy. 2007. 129 pp.

102 Nordberg, Mikael. Allies Yet Rivals: Input Joint Ventures and Their Competitive Effects. 2007. 122 pp.

103 Johansson, Fredrik. Essays on Measurement Error and Nonresponse. 2007. 130 pp.

104 Haraldsson, Mattias. Essays on Transport Economics. 2007. 104 pp.

105 Edmark, Karin. Strategic Interactions among Swedish Local Governments. 2007. 141 pp.

106 Oreland, Carl. Family Control in Swedish Public Companies. Implications for Firm Performance, Dividends and CEO Cash Compensation. 2007. 121 pp.

107 Andersson, Christian. Teachers and Student Outcomes: Evidence using Swedish Data. 2007. 154 pp.

108 Kjellberg, David. Expectations, Uncertainty, and Monetary Policy. 2007. 132 pp.

109 Nykvist, Jenny. Self-employment Entry and Survival - Evidence from Sweden. 2008. 94 pp.

110 Selin, Håkan. Four Empirical Essays on Responses to Income Taxation. 2008. 133 pp.

111 Lindahl, Erica. Empirical studies of public policies within the primary school and the sickness insurance. 2008. 143 pp.

112 Liang, Che-Yuan. Essays in Political Economics and Public Finance. 2008. 125 pp.

113 Elinder, Mikael. Essays on Economic Voting, Cognitive Dissonance, and Trust. 2008. 120 pp.

114 Grönqvist, Hans. Essays in Labor and Demographic Economics. 2009. 120 pp.

115 Bengtsson, Niklas. Essays in Development and Labor Economics. 2009. 93 pp.

116 Vikström, Johan. Incentives and Norms in Social Insurance: Applications, Identification and Inference. 2009. 205 pp.

117 Liu, Qian. Essays on Labor Economics: Education, Employment, and Gender. 2009. 133 pp.

118 Glans, Erik. Pension reforms and retirement behaviour. 2009. 126 pp.

119 Douhan, Robin. Development, Education and Entrepreneurship. 2009.

120 Nilsson, Peter. Essays on Social Interactions and the Long-term Effects of Early-life Conditions. 2009. 180 pp.

121 Johansson, Elly-Ann. Essays on schooling, gender, and parental leave. 2010. 131 pp.

122 Hall, Caroline. Empirical Essays on Education and Social Insurance Policies. 2010. 147 pp.

123 Enström-Öst, Cecilia. Housing policy and family formation. 2010. 98 pp.

124 Winstrand, Jakob. Essays on Valuation of Environmental Attributes. 2010. 96 pp.

125 Söderberg, Johan. Price Setting, Inflation Dynamics, and Monetary Policy. 2010. 102 pp.

126 Rickne, Johanna. Essays in Development, Institutions and Gender. 2011. 138 pp.

127 Hensvik, Lena. The effects of markets, managers and peers on worker outcomes. 2011. 179 pp.

128 Lundqvist, Heléne. Empirical Essays in Political and Public. 2011. 157 pp.

129 Bastani, Spencer. Essays on the Economics of Income Taxation. 2012. 257 pp.

130 Corbo, Vesna. Monetary Policy, Trade Dynamics, and Labor Markets in Open Economies. 2012. 262 pp.

131 Nordin, Mattias. Information, Voting Behavior and Electoral Accountability. 2012. 187 pp.

132 Vikman, Ulrika. Benefits or Work? Social Programs and Labor Supply. 2013. 161 pp.

133 Ek, Susanne. Essays on unemployment insurance design. 2013. 136 pp.

134 Österholm, Göran. Essays on Managerial Compensation. 2013. 143 pp.

135 Adermon, Adrian. Essays on the transmission of human capital and the impact of technological change. 2013. 138 pp.

136 Kolsrud, Jonas. Insuring Against Unemployment 2013. 140 pp.

137 Hanspers, Kajsa. Essays on Welfare Dependency and the Privatization of Welfare Services. 2013. 208 pp.

138 Persson, Anna. Activation Programs, Benefit Take-Up, and Labor Market Attachment. 2013. 164 pp.

139 Engdahl, Mattias. International Mobility and the Labor Market. 2013. 216 pp.

140 Krzysztof Karbownik. Essays in education and family economics. 2013. 182 pp.

141 Oscar Erixson. Economic Decisions and Social Norms in Life and Death Situations. 2013. 183 pp.

142 Pia Fromlet. Essays on Inflation Targeting and Export Price Dynamics. 2013. 145 pp.

143 Daniel Avdic. Microeconometric Analyses of Individual Behavior in Public Welfare Systems. Applications in Health and Education Economics. 2014. 176 pp.

144 Arizo Karimi. Impacts of Policies, Peers and Parenthood on Labor Market Outcomes. 2014. 221 pp.

145 Karolina Stadin. Employment Dynamics. 2014. 134 pp.

146 Haishan Yu. Essays on Environmental and Energy Economics. 132 pp.

147 Martin Nilsson. Essays on Health Shocks and Social Insurance. 139 pp.

148 Tove Eliasson. Empirical Essays on Wage Setting and Immigrant Labor Market Opportunities. 2014. 144 pp.

149 Erik Spector. Financial Frictions and Firm Dynamics. 2014. 129 pp.

150 Michihito Ando. Essays on the Evaluation of Public Policies. 2015. 193 pp.

151 Selva Bahar Baziki. Firms, International Competition, and the Labor Market. 2015. 183 pp.

152 Fredrik Sävje. What would have happened? Four essays investigating causality. 2015. 229 pp.

153 Ina Blind. Essays on Urban Economics. 2015. 197 pp.

154 Jonas Poulsen. Essays on Development and Politics in Sub-Saharan Africa. 2015. 240 pp.

155 Lovisa Persson. Essays on Politics, Fiscal Institutions, and Public Finance. 2015. 137 pp.

156 Gabriella Chirico Willstedt. Demand, Competition and Redistribution in Swedish Dental Care. 2015. 119 pp.

157 Yuwei Zhao de Gosson de Varennes. Benefit Design, Retirement Decisions and Welfare Within and Across Generations in Defined Contribution Pension Schemes. 2016. 148 pp.

158 Johannes Hagen. Essays on Pensions, Retirement and Tax Evasion. 2016. 195 pp.

159 Rachatar Nilavongse. Housing, Banking and the Macro Economy. 2016. 156 pp.

160 Linna Martén. Essays on Politics, Law, and Economics. 2016. 150 pp.

161 Olof Rosenqvist. Essays on Determinants of Individual Performance and Labor Market Outcomes. 2016. 151 pp.

162 Linuz Aggeborn. Essays on Politics and Health Economics. 2016. 203 pp.

163 Glenn Mickelsson. DSGE Model Estimation and Labor Market Dynamics. 2016. 166 pp.

164 Sebastian Axbard. Crime, Corruption and Development. 2016. 150 pp.

165 Mattias Öhman. Essays on Cognitive Development and Medical Care. 2016. 181 pp.

166 Jon Frank. Essays on Corporate Finance and Asset Pricing. 2017. 160 pp.

167 Ylva Moberg. Gender, Incentives, and the Division of Labor. 2017. 220 pp.

168 Sebastian Escobar. Essays on inheritance, small businesses and energy consumption. 2017. 194 pp.

169 Evelina Björkegren. Family, Neighborhoods, and Health. 2017. 226 pp.

170 Jenny Jans. Causes and Consequences of Early-life Conditions. Alcohol, Pollution and Parental Leave Policies. 2017. 209 pp.

171 Josefine Andersson. Insurances against job loss and disability. Private and public interventions and their effects on job search and labor supply. 2017. 175 pp.

172 Jacob Lundberg. Essays on Income Taxation and Wealth Inequality. 2017. 173 pp.

173 Anna Norén. Caring, Sharing, and Childbearing. Essays on Labor Supply, Infant Health, and Family Policies. 2017. 206 pp.

174 Irina Andone. Exchange Rates, Exports, Inflation, and International Monetary Cooperation. 2018. 174 pp.

175 Henrik Andersson. Immigration and the Neighborhood. Essays on the Causes and Consequences of International Migration. 2018. 181 pp.

176 Aino-Maija Aalto. Incentives and Inequalities in Family and Working Life. 2018. 131 pp.

177 Gunnar Brandén. Understanding Intergenerational Mobility. Inequality, Student Aid and Nature-Nurture Interactions. 2018. 125 pp.

178 Mohammad H. Sepahvand. Essays on Risk Attitudes in Sub-Saharan Africa. 2019. 215 pp.

179 Mathias von Buxhoeveden. Partial and General Equilibrium Effects of Unemployment Insurance. Identification, Estimation and Inference. 2019. 89 pp.

180 Stefano Lombardi. Essays on Event History Analysis and the Effects of Social Programs on Individuals and Firms. 2019. 150 pp.

181 Arnaldur Stefansson. Essays in Public Finance and Behavioral Economics. 2019. 191 pp.

182 Cristina Bratu. Immigration: Policies, Mobility and Integration. 2019. 173 pp.

183 Tamás Vasi. Banks, Shocks and Monetary Policy. 2020. 148 pp.

184 Jonas Cederlöf. Job Loss: Consequences and Labor Market Policy. 2020. 213 pp.

185 Dmytro Stoyko. Expectations, Financial Markets and Monetary Policy. 2020. 153 pp.

186 Paula Roth. Essays on Inequality, Insolvency and Innovation. 2020. 191 pp.

187 Fredrik Hansson. Consequences of Poor Housing, Essays on Urban and Health Economics. 2020. 143 pp.

188 Maria Olsson. Essays on Macroeconomics: Wage Rigidity and Aggregate Fluctuations. 2020. 130 pp.

189 Dagmar Müller. Social Networks and the School-to-Work Transition. 2020. 146 pp.

190 Maria Sandström. Essays on Savings and Intangible Capital. 2020. 129 pp.

191. Anna Thoresson. Wages and Their Impact on Individuals, Households and Firms. 2020. 220 pp.

192. Jonas Klarin. Empirical Essays in Public and Political Economics. 2020. 129 pp.

193. André Reslow. Electoral Incentives and Information Content in Macroeconomic Forecasts. 2021. 184 pp.

194. Davide Cipullo. Political Careers, Government Stability, and Electoral Cycles. 2021. 308 pp.

195. Olle Hammar. The Mystery of Inequality: Essays on Culture, Development, and Distributions. 2021. 210 pp.

196. J. Lucas Tilley. Inputs and Incentives in Education. 2021. 184 pp.

197. Sebastian Jävervall. Corruption, Distortions and Development. 2021. 215 pp.

198. Vivika Halapuu. Upper Secondary Education: Access, Choices and Graduation. 2021. 141 pp.

199. Charlotte Paulie. Essays on the Distribution of Production, Prices and Wealth. 2021. 141 pp.

200. Kerstin Westergren. Essays on Inflation Expectations, Monetary Policy and Tax Reform. 2021. 124 pp.

201. Melinda Süveg. Finance, Shocks, Competition and Price Setting. 2021. 137 pp.

202. Adrian Poignant. Gold, Coal and Iron. Essays on Industrialization and Economic Development. 2022. 214 pp.

203. Daniel Bougt. A Sequence of Essays on Sequences of Auctions. 2022. 188 pp.

204. Lillit Ottosson. From Welfare to Work. Financial Incentives, Active Labor Market Policies, and Integration Programs. 2022. 219 pp.

205. Yaroslav Yakymovych. Workers and Occupations in a Changing Labour Market. The Heterogeneous Effects of Mass Layoffs and Social Safety Nets. 2022. 212 pp.

206. Raoul van Maarseveen. Urbanization and Education. The Effect of Childhood Urban Residency on Educational Attainment. 2022. 210 pp.

207. Sofia Hernnäs. Automation and the Consequences of Occupational Decline. 2022. 229 pp.

208. Simon Ek, Structural Change, Match Quality, and Integration in the Labor Market. 2023. 176 pp.

209. Maximiliano Sosa Andres, A Risky and Polarized World: Essays on Uncertainty, Ideology and Foreign Policy. 2023. 226 pp.

210. Alice Hallman, Hypocrites, Devil's Advocates, and Bandwagoneers. Essays on Costly Signaling. 2023. 176 pp.

211. Edvin Hertegård, Essays on Families, Health Policy, and the Determinants of Children's Long-Term Outcomes. 2023. 236 pp.

212. Josefin Videnord, Trade, Innovation, and Gender. 2023. 156 pp.

213. Ao Fei, Essays on Taxation, Externalities, and Poverty Traps. 2023. 127 pp.

214. Anton Sundberg, Essays in Labor Economics. Parenthood, Immigration, and Education. 2024. 209 pp.

215. Gabriella Kindström, Urban Dynamics and Contemporary Challenges. Essays on Housing and Neighborhood Amenities. 2024. 161 pp.

216. Akib Khan, Essays in Human Capital and Immigration. 2024. 308 pp.

217. Erika Forsberg, Labor-market inequality. Essays on the roles of families, firms, location, and criminal records. 2024. 218 pp.

218. Qingyan Peng, Technology, Nonstandard Jobs, Occupations, and Firm Performance. 2024. 110 pp.

219. Hiep Nguyen, Essays on Environmental and Monetary Policy. 2024. 171 pp.

220. Markus Ridder, Essays on Polarization and Inequality. 2024. 185 pp.

221. Zunyuan Zheng, Cancer Experience and Charitable Giving. 2024. 176 pp.

222. Malin Backman, Essays on Women in the Labor Market. Technology, Inequality, and the Future of Work. 2025. 182 pp.