

Gardberg, Malin; Heyman, Fredrik; Olsson, Martin; Tåg, Joacim

Working Paper

Pre-AI sorting, post-AI inequality: Generative AI and the gender wage gap

IFN Working Paper, No. 1534

Provided in Cooperation with:

Research Institute of Industrial Economics (IFN), Stockholm

Suggested Citation: Gardberg, Malin; Heyman, Fredrik; Olsson, Martin; Tåg, Joacim (2025) : Pre-AI sorting, post-AI inequality: Generative AI and the gender wage gap, IFN Working Paper, No. 1534, Research Institute of Industrial Economics (IFN), Stockholm

This Version is available at:

<https://hdl.handle.net/10419/331456>

Standard-Nutzungsbedingungen:

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

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

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

Terms of use:

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

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

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

IFN Working Paper No. 1534, 2025

Pre-AI Sorting, Post-AI Inequality: Generative AI and the Gender Wage Gap

Malin Gardberg Fredrik Heyman, Martin Olsson and
Joacim Tåg

Pre-AI Sorting, Post-AI Inequality: Generative AI and the Gender Wage Gap*

Malin Gardberg
IFN

Fredrik Heyman
IFN and Lund University

Martin Olsson
IFN and IFAU

Joacim Tåg
IFN and Hanken School of Economics

August 2025

Abstract

We examine how gender-based occupational sorting before the release of ChatGPT relates to predicted exposure to generative AI and its potential implications for the gender wage gap. Using Swedish administrative data, we find that women are overrepresented in occupations predicted to be more affected by generative AI. Simulations based on deviations from the 2021 occupational and wage distribution—incorporating predicted AI exposure and task complementarity—show that generative AI can widen the gender wage gap through existing patterns of occupational sorting.

Keywords: Generative AI, gender wage gap, technological change, occupational sorting, complementarity.
JEL Codes: J16, J31, O33, J24

*E-mails: malin.gardberg@ifn.se, fredrik.heyman@ifn.se, martin.olsson@ifn.se, and joacim.tag@ifn.se. We are grateful to the Marianne and Marcus Wallenberg Foundation for financial support through grant 2024.0019, as well as the Jan Wallander and Tom Hedelius Foundation and the Tore Browaldh Foundation through grant P25-0107. We thank Peter Fredriksson, Siri Isaksson, Maddalena Ronchi, and Björn Tyrefors for excellent discussions and suggestions.

1 Introduction

Generative AI is reshaping the labor market by automating cognitive tasks and transforming occupational structures. Early evidence shows that tools like ChatGPT can boost productivity in specific domains, yet their broader impact on wages and employment remains uncertain.¹ Past technological advancements such as computers and robotics have altered the demand for skills and influenced gender wage dynamics. While automation in the late 20th century often displaced men more than women (Black and Spitz-Oener, 2010; Cortes and Pan., 2019), and widened the gender wage gap (Aksoy, Özcan and Philipp, 2021), the spread of personal computers helped narrow gender wage gaps by increasing demand for cognitive skills (Beaudry and Lewis, 2014). The effects of generative AI on gender wage gaps are still unknown, but this new technology disproportionately affects occupations emphasizing communication, content creation, and analytical reasoning—fields where women are often overrepresented (Felten, Raj and Seamans, 2023). Its role in shaping future wage inequality therefore warrants attention.

In this paper, we investigate how gender-based occupational sorting of workers relates to predicted future exposure to generative AI, and what this may imply for the gender wage gap. Using rich Swedish administrative data from 2021—before the public release of ChatGPT—we analyze if women are disproportionately employed in occupations predicted to be more affected by generative AI. We then assess how these differences could in the future contribute to changes in the gender wage gap by incorporating predicted exposure to generative AI and task complementarity into a simple simulation exercise.

We have three main findings. First, we document a robust gender gap in predicted exposure to generative AI. Women make up about 20 percent of workers in the least exposed occupations and over 60 percent in the most exposed. We quantify this gap in a regression framework, gradually adding a battery of controls for worker characteristics, education, and sorting across regions, industries, and firms. While the unconditional gender gap in exposure is 10 percentage points, or 22 percent, it declines to less than 4 percentage points when all controls are included, with most of the reduction explained by industry and educational

¹See e.g., Noy and Zhang (2023); Brynjolfsson, Li and Raymond (2023); Eloundou, Manning, Mishkin and Rock (2023); Acemoglu (2024).

sorting. Still, women remain significantly more exposed to generative AI than men, both unconditionally and conditional on observables.

Second, we examine how gender differences in predicted exposure to generative AI relate to wages by estimating wage regressions that include occupational exposure to generative AI. The raw gender wage gap is -8.7 percent. Including predicted generative AI exposure as a control widens the gap by almost three-quarters to -15.4 percent, because women are more exposed to generative AI yet earn less than men. This pattern persists even after controlling for education, other worker characteristics, industry, region, sector, and firm fixed effects. When these controls are included the overall wage gap reduces to -6.9 percent, yet including generative AI exposure still increases it by 1.6 percentage points or 23 percent. These findings suggest that generative AI could have substantial impacts on the gender wage gap. If future AI adoption leads to wage declines in highly exposed occupations, this pattern implies that women may be disproportionately affected, potentially widening the gender wage gap. Conversely, if generative AI raises productivity and wages in these occupations, the gap may narrow.

Finally, we simulate how generative AI could affect the gender wage gap under reasonable assumptions about future wage effects based on the observed 2021 exposure and wage distribution and assuming that the sorting pattern across occupations remains unchanged. Our framework allows future wages to depend on the predicted occupational exposure, the degree of complementarity between AI and tasks, baseline wages, and a given wage change. In our baseline setting, assuming a modest annual 3 percent wage effect over five years, the gender wage gap would decrease by almost a fifth if AI raises wages across all occupations, but would increase by almost a fifth if AI depresses wages everywhere. Most scenarios suggest that even modest negative wage effects in a small share of occupations would lead to an increase in the gender wage gap. Robustness checks varying the exposure-wage relationship and the assumed wage effects confirm that the main conclusion that generative AI may widen the gender wage gap is stable.

Our paper contributes to a growing literature that examines the heterogeneous labor market effects of automation and AI technologies, see Black and Spitz-Oener (2010); Beaudry and Lewis (2014); Cortes and Pan. (2019); Aksoy et al. (2021). While much of this literature focuses on broad employment and wage

outcomes, little is known about how generative AI, which targets high-level cognitive and language-based tasks, may interact with existing gender wage gaps. In particular, research on gender inequality emphasizes the role of occupational sorting in contributing to wage gaps (Bertrand, 2011; Goldin, 2014; Olivetti and Petrongolo, 2017; Blau and Kahn, 2017; Sin, Stillman and Fabling, 2022). These persistent sorting patterns—shaped by education, preferences, discrimination, and institutional constraints—suggest that exposure to new technologies like generative AI is unlikely to be gender-neutral. Understanding how predicted generative AI exposure maps into pre-existing occupational structures is therefore key to anticipating its future impact on gender wage gaps.

Recent evidence shows that women are less likely than men to adopt generative-AI tools (Carvajal, Franco and Isaksson, 2024; Humlum and Vestergaard, 2024; Otis, Cranney, Delecourt and Koning, 2024), even though high-performing women may stand to gain the most from using the technology (Franco, Irmert and Isaksson, 2025). Consequently, even when generative AI complements tasks and raises wages, gender gaps in adoption patterns could magnify overall wage disparities. Our mechanism also echoes Dobbie, Arnold and Hull (2024), who show that algorithms trained on selective data can amplify pre-existing group inequalities; in our setting, gender-biased occupational sorting plays a similar role, providing a structural channel through which generative AI wage shocks may widen the gender wage gap.

The value added we bring to this literature is to offer a large-scale simulation analysis of how gender-based occupational sorting before the emergence of generative AI may shape the technology’s distributional impact with respect to the gender wage gap. From this exercise, we learn that, under plausible wage impact scenarios, generative AI adoption may increase the gender wage gap due to pre-existing sorting patterns, even in the absence of behavioral responses to the technology itself.

The remainder of the paper is organized as follows. Section 2 describes the data sources, construction of the generative AI exposure measure, and empirical methodology. Section 3 presents the results, including analyses of exposure differences and wage regressions, as well as the simulation outcomes. Section 4 concludes.

2 Data

2.1 Data and sample selection

Our analysis is based on high-quality administrative micro-data from Statistics Sweden (SCB). The main source is the Salary Structures Database (*Lönestrukturstatistiken*), an annual census of wages and occupations that covers all public-sector employees and a stratified random sample of private-sector workers in Sweden. We use sampling weights supplied by SCB to obtain economy-wide, representative estimates. Wages are reported as full-time equivalents, i.e., adjusted for hours worked and absences. We supplement the data with information on gender, age, education, and industry from Statistics Sweden’s LISA database (Longitudinal Integrated Database for Health Insurance and Labour Market Studies). We restrict the sample to employed individuals with occupations that can be matched to the predicted generative AI exposure measure discussed below. The analysis focuses on the single pre-AI year 2021, thereby avoiding any labor-market effects that may have occurred after ChatGPT’s public release in November 2022.

2.2 Measuring predicted exposure to generative AI

Our measure of predicted exposure to generative AI is based on Felten et al. (2023), that quantifies how large language modeling proficiencies (such as the ability to model, predict, or mimic human language) relates to 52 human abilities relevant to occupational job tasks – accounting for ability prevalence and importance within an occupation. The measure combines information from the Electronic Frontier Foundation and the Occupational Information Network (O*NET). Each occupation’s exposure is the weighted average of the proficiencies across the 52 O*NET abilities, with higher values indicating greater exposure. Since the original scores are defined using the U.S. Standard Occupational Classification (SOC) system, we map them to the Swedish Standard Classification of Occupations (SSYK 2012). The mapping yields 365 occupations with exposure scores that we link to individual workers.

A limitation of the predicted exposure score is that it does not tell us if exposure is expected to have a positive or negative effect on wages in a given occupation. To obtain a measure of this, we supplement the

data with the Pizzinelli, Panton, Tavares, Cazzaniga and Li (2023) measure of predicted AI complementarity, which captures the potential for AI to augment human tasks at the occupational level.² Although originally developed using a broader set of AI technologies from Felten, Raj and Seamans (2021), generative AI is an essential part of this measure. The combination of the exposure measure and the complementarity measure allows us to simulate how a hypothetical wage change could impact the gender wage gap in the economy. For ease of interpretation, we standardize both measures to range between zero and one.

Table 1 provides summary statistics on our final sample of 2,516,762 observations. On average, individuals earn 37.3 TSEK per month (full-time equivalent), have 19.4 years of labor market experience, and have 5.4 years of tenure with their current employer.³ Women earn less than men on average (35.4 vs. 39.3 TSEK), have slightly less experience (18.8 vs. 20.0 years), and have shorter tenure (5.0 vs. 5.8 years). Educational attainment differs significantly, with 37.8 percent of women having at least two years of university education compared to 24.6 percent of men. The share of immigrants is similar across genders. A higher share of women work in the public sector (45.3 percent vs. 16.4 percent), whereas men are more likely to be employed in the private sector. Women also exhibit higher exposure to generative AI (0.548 vs. 0.448) and slightly lower measured complementarity with AI technologies (0.506 vs. 0.525).

3 Analysis

3.1 Are women working in occupations predicted to be more exposed to generative AI?

We begin by examining the predicted exposure to generative AI and whether it varies by gender. In our sample, men are disproportionately employed in less exposed occupations (mean exposure 0.448), while women are more concentrated in highly exposed ones (mean exposure 0.548). Figure 1(a) illustrates this pattern: women make up roughly 20 percent of workers in the least exposed occupations, but over 60 percent in the most exposed.

Several reasons may contribute to the pattern that women are more likely than men to work in occu-

²Note that we use the explicit measure of potential complementarity of AI from Pizzinelli et al. (2023), which they label as θ .

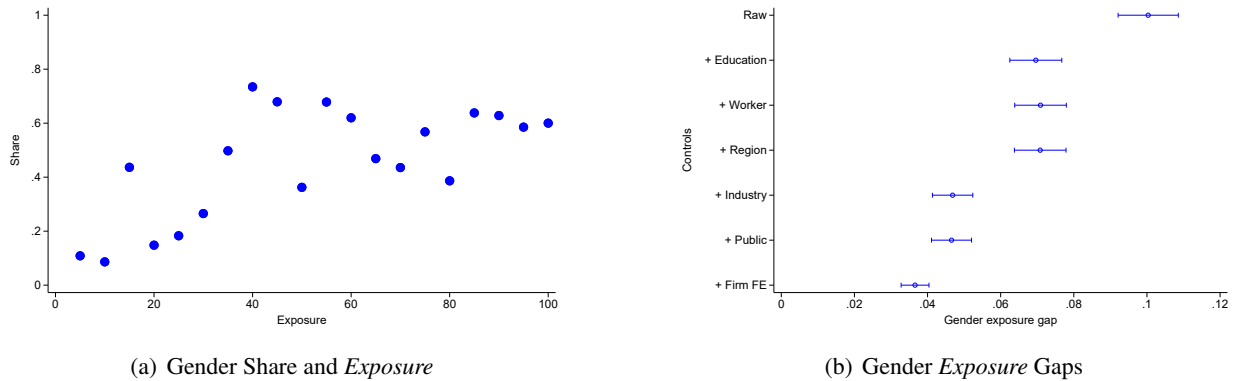
³In May 2025, the USD to SEK exchange rate was 9.6.

Table 1: Summary statistics

	All (1)	Females (2)	Males (3)	Diff. (2)-(3) (4)
Wage	37.348	35.379	39.326	-3.947
Experience	19.398	18.788	20.012	-1.224
Tenure	5.383	4.971	5.797	-0.826
Education	0.312	0.378	0.246	0.132
Immigrant	0.234	0.238	0.230	0.009
Public	0.309	0.453	0.164	0.289
Exposure	0.498	0.548	0.448	0.100
Complementarity	0.516	0.506	0.525	-0.019

Notes: *Wage* denotes full-time equivalent monthly wages in thousands of SEK. *Experience* is defined based on age and education: age minus 16 for individuals with only primary education or missing data; minus 19 for upper secondary; minus 20 for short post-secondary; minus 23 for long post-secondary; and age minus graduation year for higher education. *Tenure* is measured using worker-firm links from 1990 to 2021. *Education* equals one if the individual has at least two years of university education, and zero otherwise. *Immigrant* is an indicator of being an immigrant. *Public* is an indicator of public sector employment. *Exposure* refers to the generative AI exposure measure (AIOELM) in Felten et al. (2023), while *Complementarity* captures the potential for AI task augmentation (theta) from Pizzinelli et al. (2023). All differences reported in Column 4, except for *Immigrant*, are statistically significant at the one percent level. The sample consists of 2,516,762 observations that are weighted with sampling weights from SCB to form an economy-wide representative sample for the year 2021.

Figure 1: The Gender Gap in *Exposure*



Notes: Panel (a) plots the share of women across the *Exposure* distribution (divided into vigintiles), where higher values reflect greater exposure. Panel (b) reports the estimated gender gap in *Exposure* (i.e., the coefficient β from Eq. (1)) with 95% confidence intervals, adding controls stepwise. *Worker* controls include *Experience*, *Tenure*, and *Immigrant*.

pations with higher exposure to generative AI. For instance, women may be disproportionately employed in administrative, communication, legal, educational, and other service-oriented occupations that involve routine cognitive and information-processing tasks, which are areas that generative AI is particularly well suited to automate or augment work. Occupational sorting by gender, shaped by historical norms, educational choices, and labor market discrimination, may also play a role.

To explore more in detail why women are more likely than men to work in occupations with higher exposure to generative AI, we estimate the following model at the worker level:

$$\text{Exposure}_{o(i)} = \alpha + \beta \text{Female}_i + X_i' \theta + \varepsilon_i \quad (1)$$

where $\text{Exposure}_{o(i)}$ denotes the predicted exposure to generative AI for occupation o , and Female_i is an indicator equal to one if individual i is a woman. The vector X_i includes worker-level fixed effects for experience, tenure, and immigrant status, as well as fixed effects for public employment, regions (24 counties), industries (16 industries), and firms (39,267). We add controls stepwise to assess their explanatory power. The coefficient β captures the average gender gap in exposure, conditional on observed characteristics and sorting patterns.

The unconditional gap in Figure 1(b) shows that women are, on average, 10 percentage points, or 22 percent, more exposed to generative AI than men when the model does not account for any observable differences. When controlling for education – where 38 percent of women have a university education compared to 25 percent of men in our sample, the exposure gap decreases to just under 7 percentage points.⁴ Thus, despite controlling for women being overrepresented among university-educated individuals, they remain significantly (both statistically and economically) more exposed to generative AI than men. The gap remains nearly unchanged when other worker characteristics and regional fixed effects are added to the model. However, it narrows to approximately 5 percentage points when industry and sector fixed effects are included. Finally, when accounting for firm-level sorting over and above the other control variables, the estimated gap diminishes to 3.7 percentage points. Thus, sorting across industries, sectors, and especially

⁴Unreported results show that exposure is higher among more educated workers.

firms explains roughly 60 percent of the initial difference, yet a sizeable and statistically significant exposure gap in favor of women remains.

In summary, the results in Figure 1 show that women tend to be working in occupations that are predicted to be substantially more exposed to generative AI than men.

3.2 Could women’s higher predicted exposure to generative AI influence the gender wage gap?

We next investigate whether women’s over-representation in occupations that are predicted to be more exposed to generative AI technologies contributes to the gender wage gap. To do so, we estimate the gender wage gap in two models—one that excludes the predicted occupational exposure to generative AI and one that includes it. Comparing the estimated gender-gap coefficients from these models tells us whether the observed pattern of female sorting into more AI-exposed occupations was already associated with systematic wage differentials between men and women in 2021. This could reflect, for instance, higher educational or skill requirements in those occupations. If such a correlation holds, future wage shocks aligned with the occupational exposure profile of generative AI could feed through to the gender wage gap.

Formally, we estimate the following regression:

$$\log(w_i) = \alpha + \beta \text{Female}_i + \gamma \text{Exposure}_{o(i)} + X_i' \theta + \varepsilon_i, \quad (2)$$

where w_i is the full-time equivalent monthly wage for individual i , and the coefficient β captures the average log wage gap between women and men, controlling for occupational exposure ($\text{Exposure}_{o(i)}$) and X_i . The vector X_i includes *Education*, *Experience*, *Tenure*, *Immigrant*, *Public*, and fixed effects for regions, industries, and firms. We cluster the standard errors at the firm level.

Table 2 presents the estimation results. Panel A excludes predicted occupational AI exposure as a control, while Panel B includes it. In the unconditional specification (Column 1, Panel A), the raw gender wage gap is -8.7 percent. When predicted occupational AI exposure is introduced as the sole control (Column 1, Panel B), the wage gap increases substantially to -15.4 percent. This 6.7 percentage point

Table 2: The Gender Wage Gap and *Exposure*

Panel A: Not controlling for <i>Exposure</i>					
	(1)	(2)	(3)	(4)	(5)
<i>Female</i>	-0.087*** (0.005)	-0.126*** (0.004)	-0.116*** (0.003)	-0.076*** (0.003)	-0.068*** (0.003)
R^2	0.02	0.18	0.29	0.36	0.55
Panel B: Controlling for <i>Exposure</i>					
	(1)	(2)	(3)	(4)	(5)
<i>Female</i>	-0.154*** (0.004)	-0.162*** (0.004)	-0.150*** (0.003)	-0.099*** (0.003)	-0.085*** (0.003)
<i>Exposure</i>	0.667*** (0.011)	0.522*** (0.012)	0.482*** (0.011)	0.483*** (0.010)	0.469*** (0.007)
R^2	0.25	0.30	0.39	0.44	0.60
Controls					
<i>Education</i>	No	Yes	Yes	Yes	Yes
<i>Worker</i>	No	No	Yes	Yes	Yes
<i>Industry</i>	No	No	No	Yes	Yes
<i>Region and Public</i>	No	No	No	No	Yes
<i>Firm</i>	No	No	No	No	Yes
N	2,516,762	2,516,762	2,516,762	2,516,762	2,516,762

Notes: This table presents estimates of the gender log wage gap using Eq. (2) with and without the *Exposure* as a control. *Worker* controls include *Education*, *Experience*, *Tenure*, and *Immigrant*.

increase—equivalent to a 73 percent rise—suggests that women tend to sort into occupations with higher predicted AI exposure, which, in 2021, were also associated with higher average wages. This pattern persists even after controlling for additional observable characteristics. In the most saturated specification (Column 5), the inclusion of predicted occupational AI exposure continues to widen the wage gap to –8.5 percent—an increase of 1.6 percentage points, or approximately 24 percent.

Overall, these results indicate that women are overrepresented in occupations with higher predicted exposure to generative AI—occupations that, as of 2021, were already associated with higher average wages. If future adoption of generative AI leads to wage declines in these highly exposed occupations, this pattern suggests that women could be disproportionately adversely affected, potentially exacerbating the gender wage gap. Conversely, if AI adoption enhances productivity and leads to wage gains in these roles, the gap may narrow, offering a potential equalizing effect.

3.3 Simulating effects on the gender wage gap

As shown above, women are over-represented in occupations with higher predicted exposure to generative AI technologies, and this sorting is associated with the 2021 gender wage gap—even before the public release of ChatGPT could have had any impact on wages. We now simulate the potential future impact of generative AI on the gender wage gap, based on plausible assumptions about how wages may evolve across the 2021 exposure distribution. These simulations provide a framework for assessing how specific deviations from the observed exposure–wage patterns in 2021 could influence the gender wage gap going forward.

The direction and magnitude of generative AI’s impact on the gender wage gap will depend critically on two factors: (i) how strongly an occupation is expected to be affected by generative AI, and (ii) whether generative AI translates into wage gains or losses at the occupational level, i.e., whether it is expected to complement (+) or substitute (–) the occupation’s tasks. Our simulation framework is designed to account

for both of these dimensions. Specifically, we assume that

$$\hat{w}_i = f(Exposure_{o(i)}, Complementarity_{o(i)}, w_i, \delta), \quad (3)$$

where \hat{w}_i denotes the simulated future wage of individual i , $Exposure_{o(i)}$ is the occupational exposure measure by Felten et al. (2023), $Complementarity_{o(i)}$ is the occupational complementarity measure by Pizzinelli et al. (2023), w_i is the observed individual wage in 2021, and δ represents the assumed wage effect attributable to generative AI. This function captures the idea that future wage outcomes depend on the occupation's predicted exposure to AI, its complementarity with AI-driven tasks, baseline wages, and an assumed upper bound on the potential impact of generative AI on occupational wages.

In our baseline linear predicted exposure specification, we take the function $f(\cdot)$ to be specified as:

$$\hat{w}_i = \begin{cases} w_i(1 - Exposure_{o(i)}\delta) & \text{if } Complementarity_{o(i)} \leq d \\ w_i(1 + Exposure_{o(i)}\delta) & \text{if } Complementarity_{o(i)} > d, \end{cases} \quad (4)$$

where d represents a cutoff value in the complementarity distribution, varied over its deciles. Occupations below the complementarity cutoff are assumed to experience exposure-scaled wage losses (substitution), whereas those above the cutoff enjoy exposure-scaled gains (complementarity), as $Exposure$ ranges between 0 and 1. We further assume a cumulated 3% annual wage effect over five years ($\delta = 1.03^5 \approx 0.159$). The idea here is to capture the notion that future wage changes depend on current wages and an assumed wage change that accounts for both $Exposure$ and $Complementarity$.

To translate these generative AI-induced counterfactual wages into an economy-wide measure of the gender wage gap, we estimate the following regression model:

$$\log(w_i) - \log(\hat{w}_i) = \alpha_s + \beta_s Female_i + \varepsilon_i, \quad (5)$$

where the dependent variable captures the log difference between observed and simulated wages, and β_s measures the simulated gender wage gap. More specifically, β_s measures the change in the gender wage

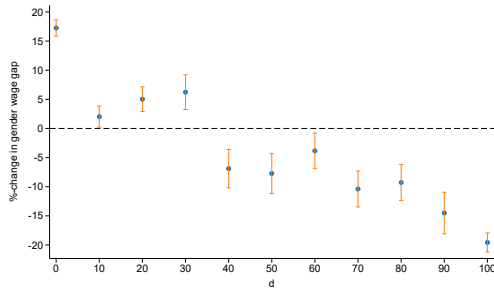
gap implied by a given (d, δ) pair. The results are presented in Figure 2, which plots estimates of β_s across different cutoff values d in the complementarity distribution (measured in deciles). Positive β_s values indicate a narrowing of the gap relative to its 2021 baseline; negative values imply a widening.

Figure 2(a) shows results from the baseline linear exposure simulation introduced above. Consider the first estimate at $d = 0$, which assumes that generative AI raises wages in all occupations, and where the wage gain is scaled by the exposure to generative AI (i.e., the individuals with the highest exposure to generative AI experience a 3% annual wage gain, and those that are least exposed experience no change). Under this scenario, the gender wage gap decreases by approximately 17%. The second estimate, at $d = 10$, assumes that only occupations in the bottom decile of the complementarity distribution suffer wage losses proportional to their generative AI exposure, while all others receive exposure-scaled wage gains. In this case, the gender wage gap narrows by about 2%. At $d = 50$, where all the individuals in the lower half of the complementarity distribution suffer exposure-scaled wage losses, and the individuals in occupations in the upper half of the complementarity distribution enjoy exposure-scaled wage gains, the gender wage gap is estimated to increase by around 7%. The final estimate, at $d = 100$, shows that the gender wage gap would increase by approximately 20% if generative AI were to cause wage declines across all occupations.

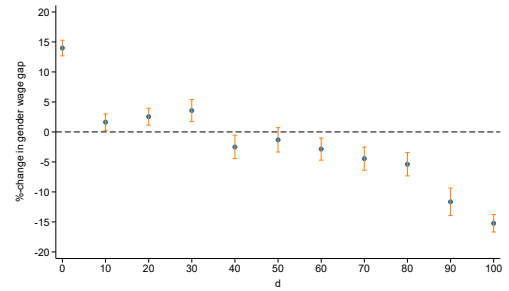
Taken together, the results across all values of d reinforce a central intuition: because women are over-represented in highly AI-exposed jobs, uniform wage gains from AI adoption tend to compress the gender wage gap, whereas uniform losses would tend to widen it. The figure yields two main conclusions. First, under a modest assumed five-year impact of generative AI on wages, the overall effect on the gender wage gap can be considerable, ranging from a decrease of about 17 percent to an increase of about 20 percent. Second, the gap expands for most complementarity cut-offs, implying that even if wage losses are confined to a small subset of occupations, the overall gender wage gap is still likely to widen.

Figure 2(a) shows the effect on the gender wage gap when keeping the interaction between exposure and wages constant while varying how this interaction changes across the complementarity distribution. As a robustness check, we first run three additional simulations to verify that our findings are not driven by the linearity assumption in Equation (4). These specifications alter the way $Exposure_{o(i)}$ interacts with the scal-

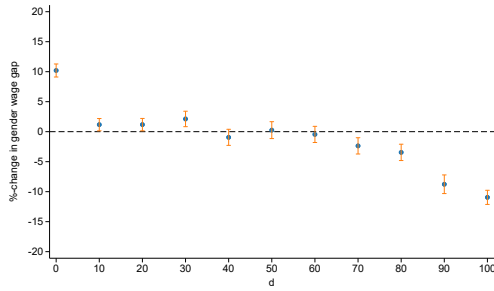
Figure 2: Simulating Changes in the Gender Wage Gap



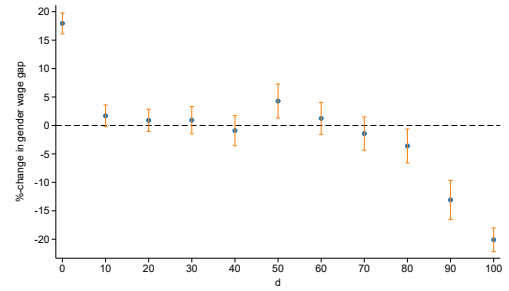
(a) Linear effect across *Exposure*



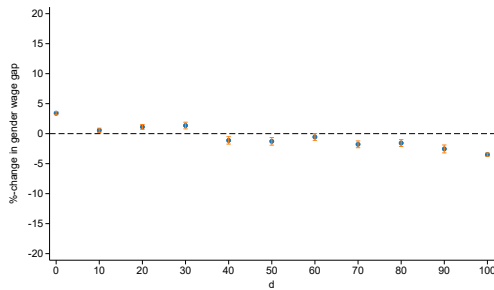
(b) Quadratic effect across *Exposure*



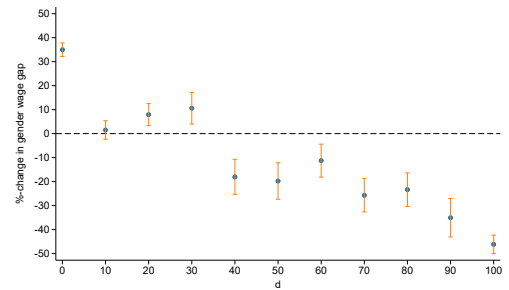
(c) Cubic effect across *Exposure*



(d) Zero effect below median and linear above across *Exposure*



(e) Short-run linear effect across *Exposure*



(f) Long-run linear effect across *Exposure*

Notes: This figure displays the simulated change in the gender wage gap (y-axis). Positive β_s values indicate a narrowing of the gap relative to its 2021 baseline; negative values imply a widening. The x-axis displays the cutoff point in the distribution of *Complementarity_o* where the wage effect changes from positive to negative. Panels (a)-(d) assume a 3% annual wage change over 5 years in the most exposed occupation in our sample in 2021, while panel (e) assumes a 3% annual wage change over 1 year, and panel (f) 3% annual wage change over 10 years.

ing parameter δ . Figures 2(b)–(d) present these results. Specifically, Figure 2(b) replaces $Exposure_{o(i)}$ with $Exposure_{o(i)}^2$ in Equation 4, Figure 2(c) replaces $Exposure_{o(i)}$ with $Exposure_{o(i)}^3$, and Figure 2(d) replaces $Exposure_{o(i)}$ with $IExposure_{o(i)}$, where $I = 1$ at the median of $Complementary_{o(i)}$ and $I = 0$ otherwise. We also illustrate what the results look like at different time horizons, where panel (e) displays the effect of a 3 percent annual change after 1 year, and panel (f) the impact of a 3 percent annual change after 10 years, while keeping the exposure term linear. Overall, Figures 2(b)–(f) illustrate that the general insights from Figure 2(a) continue to hold: even under alternative, non-linear exposure mappings or different wage shock magnitudes, the simulated effect of generative AI on the gender wage gap tend to be considerable and, if anything, tends to widen rather than narrow the gender wage gap.

4 Concluding Remarks

This paper examines how gender-based occupational sorting prior to the emergence of generative AI relates to predicted exposure to the technology and its implications for the gender wage gap. Using comprehensive Swedish administrative data from 2021, we show that women are disproportionately employed in occupations predicted to be more affected by generative AI. We further simulate potential wage changes under different assumptions on deviations from the 2021 sorting and wage distributions across occupations. These simulations indicate that generative AI can impact the gender wage gap and that even modest negative wage effects in a small share of occupations would lead to an increase in the gender wage gap. This finding is robust across specifications and assumptions.

Our findings contribute to a growing literature on the distributional consequences of technological change by highlighting the role of pre-existing labor market structures for future inequality. As generative AI continues to evolve and diffuse, understanding how exposure interacts with occupational sorting and wage structures is important for anticipating the effect on the gender wage gap.

References

Acemoglu, Daron, “The Simple Macroeconomics of AI,” NBER Working Papers 32487, National Bureau of Economic Research, Inc May 2024.

Aksoy, Cevat Giray, Berkay Özcan, and Julia Philipp, “Robots and the gender pay gap in Europe,” *European Economic Review*, 2021, 134.

Beaudry, Paul and Ethan Lewis, “Do Male-Female Wage Differentials Reflect Differences in the Return to Skill? Cross-City Evidence from 1980-2000,” *American Economic Journal: Applied Economics*, April 2014, 6 (2), 178–94.

Bertrand, Marianne, “New Perspectives on Gender,” in O. Ashenfelter and D. Card, eds., *Handbook of Labor Economics*, Vol. 4 of *Handbook of Labor Economics*, Elsevier, 2011, chapter 17, pp. 1543–1590.

Black, Sandra E. and Alexandra Spitz-Oener, “Explaining Women’s Success: Technological Change and the Skill Content of Women’s Work,” *The Review of Economics and Statistics*, February 2010, 92 (1), 187–194.

Blau, Francine D. and Lawrence M. Kahn, “The Gender Wage Gap: Extent, Trends, and Explanations,” *Journal of Economic Literature*, September 2017, 55 (3), 789–865.

Brynjolfsson, Erik, Danielle Li, and Lindsey Raymond, “Generative AI at Work,” Technical Report, National Bureau of Economic Research, Cambridge, MA April 2023.

Carvajal, Daniel, Catalina Franco, and Siri Isaksson, “Will Artificial Intelligence Get in the Way of Achieving Gender Equality?,” 2024.

Cortes, Patricia and Jessica Pan., “Gender, Occupational Segregation, and Automation,” Technical Report, Economics Studies at Brookings 2019.

- Dobbie, Will, David Arnold, and Peter Hull**, “Building Non-Discriminatory Algorithms in Selected Data,” *American Economic Review: Insights*, 2024. Forthcoming.
- Eloundou, Tyna, Sam Manning, Pamela Mishkin, and Daniel Rock**, “GPTs Are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models,” *arXiv preprint arXiv:2303.10130*, 2023.
- Felten, Edward, Manav Raj, and Robert Seamans**, “Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses,” *Strategic Management Journal*, 2021, 42 (12), 2195–2217.
- Felten, Edward W, Manav Raj, and Robert Seamans**, “Occupational Heterogeneity in Exposure to Generative AI,” *SSRN Electronic Journal*, 2023.
- Franco, Catalina, Natalie Irmert, and Siri Isaksson**, “Will AI Help or Hurt Learning?,” 2025. Mimeo.
- Goldin, Claudia**, “A Grand Gender Convergence: Its Last Chapter,” *American Economic Review*, 2014, 104 (4), 1091–1119.
- Humlum, Anders and Emilie Vestergaard**, “The Adoption of ChatGPT,” *SSRN Electronic Journal*, 2024.
- Noy, Shakked and Whitney Zhang**, “Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence,” *Science*, July 2023, 381 (6654), 187–192.
- Olivetti, Claudia and Barbara Petrongolo**, “The Economic Consequences of Family Policies: Lessons from a Century of Legislation in High-income Countries,” *Journal of Economic Perspectives*, 2017, 31 (1), 205–230.
- Otis, Nicholas G., Katelyn Cranney, Solene Delecourt, and Rembrand Koning**, “Global Evidence on Gender Gaps and Generative AI,” OSF Preprints h6a7c, Center for Open Science October 2024.
- Pizzinelli, Carlo, Augustus J Panton, Marina Mendes Tavares, Mauro Cazzaniga, and Longji Li**, “Labor Market Exposure to AI: Cross-country Differences and Distributional Implications,” *IMF Working Papers*, 2023, (23/216).

Sin, Isabelle, Steven Stillman, and Richard Fabling, “What Drives the Gender Wage Gap? Examining the Roles of Sorting, Productivity Differences, Bargaining, and Discrimination,” *The Review of Economics and Statistics*, 07 2022, 104 (4), 636–651.