


SHORT PAPER OPEN ACCESS

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

Malin Gardberg¹ | Fredrik Heyman¹ | Martin Olsson^{1,2,3} | Joacim Tåg^{1,4} 

¹IFN - Research Institute of Industrial Economics, Stockholm, Sweden | ²IFAU, Uppsala, Sweden | ³RF Berlin, Berlin, Germany | ⁴Hanken School of Economics, Helsinki, Finland

Correspondence: Joacim Tåg (joacim.tag@ifn.se)

Received: 12 June 2025 | **Revised:** 15 January 2026 | **Accepted:** 26 April 2026

Keywords: complementarity | gender wage gap | generative AI | occupational sorting | technological change

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 document that women are over-represented in occupations predicted to be more affected by generative AI. Mechanical partial-equilibrium simulations, based on hypothesized deviations from the 2021 occupational and wage distribution and incorporating predicted AI exposure and task complementarity, indicate that generative AI may widen the gender wage gap through existing patterns of gender-based occupational sorting.

JEL Classification: J16, J31, O33, J24

1 | Introduction

Generative AI has the potential to reshape the labour market by automating or augmenting cognitive tasks and changing occupational structures. Early evidence shows that large language models can boost productivity in specific domains, yet their broader impact on wages and employment remains uncertain.¹ Previous waves of technological change, such as computerization and robotics, also altered the task and skill content of jobs and affected gender wage dynamics. Computerization reduced the importance and returns to routine tasks while raising the demand for analytical and interactive tasks, contributing to a narrowing of the gender wage gap as women benefited from rising returns to cognitive skills [5–7]. Automation also reshaped occupational structures as women, despite higher initial exposure to routine work, were more likely than men to reallocate toward high-skill, high-wage occupations, contributing to declining occupational segregation [8]. By contrast, industrial robots have primarily substituted for manual and routine tasks, with heterogeneous

implications for the gender wage gap across countries and industries [9, 10].

Generative AI differs from these earlier waves of automation in that it targets language, communication, and content-creation tasks that are central to non-routine cognitive work [11]. These tasks are particularly prevalent in female-intensive occupations, such as teaching, healthcare, and administrative and clerical roles. While the effects of generative AI on gender wage gaps are still unknown, the fact that these female-intensive occupations are predicted to be disproportionately exposed suggests that generative AI may be relevant for future wage inequality.

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, combined with occupation-level indices of generative AI exposure created by Felten et al. [11], we analyse

This is an open access article under the terms of the [Creative Commons Attribution-NonCommercial](https://creativecommons.org/licenses/by-nc/4.0/) License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited and is not used for commercial purposes.

© 2026 The Author(s). *Oxford Bulletin of Economics and Statistics* published by Oxford University and John Wiley & Sons Ltd.

whether women are disproportionately employed in occupations predicted to be more affected by this technology. We then assess whether these differences could contribute to future 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% of workers in the least exposed occupations but over 60% in the most exposed ones. To quantify this gap, we estimate a series of regressions that progressively incorporate controls for worker characteristics, education and sorting across regions, industries, and firms. The difference in exposure between women and men is 10 percentage points, implying that women's exposure is approximately 22% higher than men's. Once all controls are included, the gap shrinks to less than 4 percentage points, with most of the reduction driven by industry and educational sorting. Nevertheless, even conditional on a rich set of observables, women remain significantly more exposed to predicted generative AI exposure than men.

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% . Including predicted generative AI exposure as a control widens the gap by almost three-quarters, to -15.4% , because women are more likely than men to work in occupations with higher predicted AI exposure, which in 2021 were also higher-paying on average. 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 -6.8% , yet including generative AI exposure still increases it by 1.7 percentage points or 25%. These findings suggest that wage changes aligned with predicted generative AI exposure could have substantial implications for 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 mechanically simulate how hypothetical wage changes aligned with predicted generative AI exposure could affect the gender wage gap under reasonable hypotheses 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. Specifically, our simulation framework maps a static, pre-AI exposure and complementarity profile into counterfactual future wages in partial equilibrium, holding occupational sorting, labour supply, and general-equilibrium feedback effects fixed. The resulting estimates should therefore be interpreted as illustrative, mechanical calculations rather than predictions of realized wage effects or causal impacts of generative AI adoption.

In our baseline setting, assuming a modest annual 3% wage effect over 5 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 changes in a small share of occupations would imply a wider 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 conclusions furthermore remain unchanged when extending the analysis to earned wage income, adjusting for part-time work, and allowing for intensive-margin adjustments in working hours.

Our paper contributes to a growing literature that examines the heterogeneous labour market effects of automation and AI technologies [5, 6, 8, 9, 12]. 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 emphasises the role of occupational sorting in contributing to wage gaps [13–17]. 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.

In addition, recent work on occupation-level indices of generative AI exposure points out that women may be differentially exposed compared to men. Felten et al. [11] develop the exposure index we use and show which occupations and industries are most affected by generative AI. They also document that more exposed occupations tend to pay higher wages. Gmyrek, Berg and Bescond [18] extend this line of work to a global setting by using GPT-4 to construct task-level exposure scores, showing that generative AI is more likely to augment than fully automate jobs, with particularly strong exposure in female-intensive clerical occupations. Building on this literature, Pizzinelli, Panton, Tavares, Cazzaniga and Li [19] use a combined exposure-complementarity index to document that women and highly educated workers are, in most countries, more exposed to AI, but that this exposure spans occupations with both high and low expected task complementarity. Using the same framework, Cazzaniga, Panton, Li, Pizzinelli and Tavares [20] emphasise that women's higher exposure reflects their concentration in cognitive and professional occupations and entails both potential productivity gains and heightened displacement risks.²

Our contribution to this literature is to study the implications of generative AI for the gender wage gap by combining occupation-level measures of predicted AI exposure and task complementarity with comprehensive administrative worker-level wage data for an entire economy. We then use these inputs within a simple simulation framework. We do not estimate the causal effect of AI adoption on wages. Instead, we quantify how existing gender-based occupational sorting, observed prior to the diffusion of generative AI, would mechanically interact with hypothetical AI-driven wage profiles under alternative assumptions about exposure and task complementarity. By mapping static, pre-AI exposure and complementarity measures into

counterfactual wage changes in partial equilibrium, our analysis isolates the role of current occupational sorting in shaping potential distributional outcomes. Whereas existing exposure-based studies primarily document who is more exposed to AI and in some cases how exposure relates to labour-market transitions, they do not assess how different plausible AI-induced wage profiles would translate into changes in the aggregate gender wage gap.

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, and the simulation outcomes. Section 4 concludes.

2 | Data

2.1 | Data and Sample Selection

Our analysis is based on high-quality administrative microdata 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, that is, 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 labour-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. [11], which quantifies how large language modelling proficiencies (such as the ability to model, predict, or mimic human language) relate 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 US 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 et al. [19] measure of predicted AI complementarity, which captures the potential for AI to augment human tasks at the occupational level.³ Although originally developed using a broader set of AI technologies from Felten, Raj and Seamans [25], generative AI is an essential part of this measure. The combination of the exposure measure and the complementarity measure allows us to simulate how hypothetical wage changes could translate into changes in the gender wage gap. For ease of interpretation, we standardize both measures to range between zero and one.

Both the exposure and complementarity measures should be interpreted with important caveats. First, the underlying task descriptions and expert assessments were developed before the rapid diffusion of large generative models after 2021. As a result, there is uncertainty about how well these indices capture the capabilities, use cases, and economic relevance of generative AI as of 2025 and beyond, particularly as models evolve and firms adapt their production processes. Second, the occupational exposure and complementarity scores are time-invariant in our analysis and abstract from endogenous adjustments in task content, adoption intensity, and worker reallocation across occupations. Our simulation framework thus maps a static, pre-AI exposure and complementarity profile into counterfactual future wages in partial equilibrium, holding occupational sorting, labour supply, and general-equilibrium feedback effects fixed. The resulting estimates should therefore be interpreted as illustrative, mechanical calculations rather than predictions of realized wage effects or causal impacts of generative AI adoption. Despite these limitations, the measures provide a transparent way to assess how existing gender-based occupational sorting may interact with plausible AI-related wage profiles to shape future gender wage inequality.

Table 1 provides summary statistics on our final sample of 2,516,769 observations. On average, individuals earn 37.3 TSEK per month (full-time equivalent), have 19.4 years of labour market experience, and have 5.4 years of tenure with their current employer.⁴ 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% of women having at least 2 years of university education compared to 24.6% of men. The share of immigrants is similar across genders. A higher share of women work in the public sector (45.3% vs. 16.4%), 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 1a illustrates this

pattern: women make up roughly 20% of workers in the least exposed occupations but over 60% in the most exposed.

Several reasons may contribute to the pattern that women are more likely than men to work in occupations 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 labour market discrimination, may also play a role.

TABLE 1 | Summary statistics.

	All (1)	Females (2)	Males (3)	Diff. (2-3) (4)
<i>Wage</i>	37.347	35.379	39.326	-3.947
<i>Experience</i>	19.398	18.788	20.012	-1.225
<i>Tenure</i>	5.383	4.971	5.797	-0.826
<i>Education</i>	0.312	0.378	0.246	0.132
<i>Immigrant</i>	0.234	0.238	0.230	0.009
<i>Public</i>	0.309	0.453	0.164	0.289
<i>Exposure</i>	0.498	0.548	0.448	0.100
<i>Complementarity</i>	0.516	0.506	0.525	-0.019

Note: *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 2 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. [11], while *Complementarity* captures the potential for AI task augmentation (theta) from Pizzinelli et al. [19]. All differences reported in Column 4, except for *Immigrant*, are statistically significant at the 1% level. The sample consists of 2,516,769 observations that are weighted with sampling weights from SCB to form an economy-wide representative sample for the year 2021.

To explore in more 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:

$$Exposure_{o(i)} = \alpha + \beta Female_i + X_i' \theta + \epsilon_i \quad (1)$$

where $Exposure_{o(i)}$ 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 1b shows that women are, on average, 10 percentage points or 22% more exposed to generative AI than men when the model does not account for any observable differences. When controlling for education (where 38% of women have a university education compared to 25% of men in our sample), the exposure gap decreases to just under 7 percentage points.⁵ 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 firms explains roughly 60% of the initial difference, yet a sizeable and statistically significant exposure gap in favour 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.

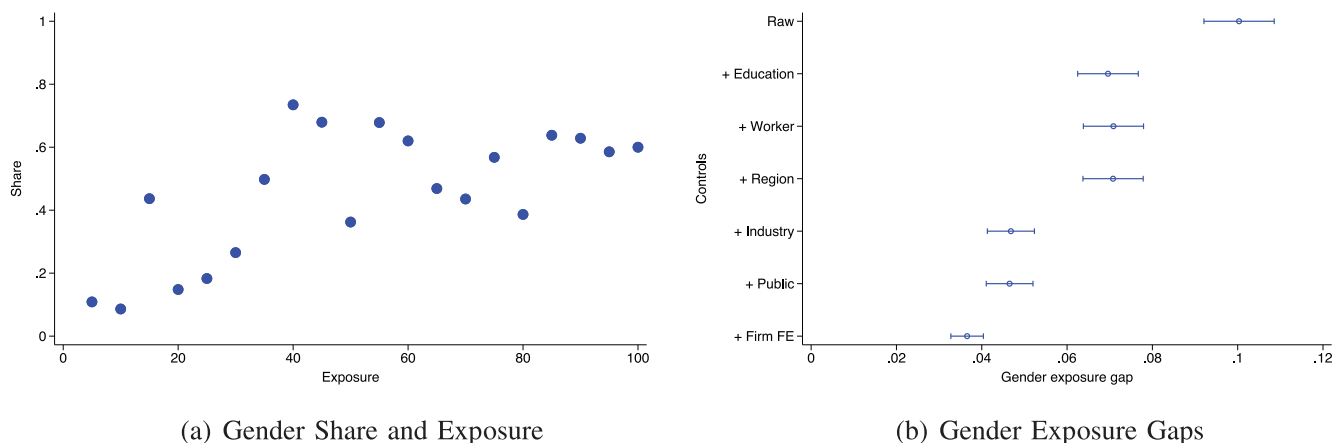


FIGURE 1 | The gender gap in *Exposure*. Panel (a) plots the share of women across the *Exposure* distribution (divided into ventiles), where higher values reflect greater exposure. Panel (b) reports the estimated gender gap in *Exposure* (i.e., the coefficient β from Equation (1)) with 95% confidence intervals, adding controls stepwise. *Worker* controls include *Experience*, *Tenure*, and *Immigrant*. (a) Gender share and *Exposure*. (b) Gender exposure Gaps. [Colour figure can be viewed at [wileyonlinelibrary.com](https://onlinelibrary.wiley.com)]

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)
<i>R</i> 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*** (0.012)	0.483*** (0.010)	0.468*** (0.007)
<i>R</i> ²	0.25	0.30	0.39	0.44	0.60
<i>Controls</i>					
<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
<i>N</i>	2,516,769	2,516,769	2,516,769	2,516,769	2,516,769

Note: This table presents estimates of the gender log wage gap using Equation (2) with and without the *Exposure* as a control. *Worker* controls include *Education*, *Experience*, *Tenure* and *Immigrant*. *** indicates statistical significance at the 1% level.

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 is related 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:

$$\log(w_i) = \alpha + \beta Female_i + \gamma Exposure_{o(i)} + X_i' \theta + \varepsilon_i \quad (2)$$

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 ($Exposure_{o(i)}$) 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. It is important to note that Equation (2) is descriptive rather than causal, relating wages in 2021 to predicted occupational exposure measured before the diffusion of generative AI.

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%. When predicted occupational AI exposure is introduced as the sole control (Column 1, Panel B), the wage gap increases substantially to −15.4%. This 6.7 percentage point increase, equivalent to a 77% rise, reflects sorting of women into occupations with higher predicted AI exposure, which in 2021 were also higher-paying on average. 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%, an increase of 1.7 percentage points, or approximately 25%.

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 Changes in 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 potential future changes in 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 the simulated change in the gender wage gap 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, that is, 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

$$\hat{w}_i = f(\text{Exposure}_{o(i)}, \text{Complementarity}_{o(i)}, w_i, \delta) \quad (3)$$

where \hat{w}_i denotes the simulated future wage of individual i , $\text{Exposure}_{o(i)}$ is the occupational exposure measure by Felten et al. [11], $\text{Complementarity}_{o(i)}$ is the occupational complementarity measure by Pizzinelli et al. [19], 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 potential occupational wage changes.

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

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

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 5 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 analyse how these generative AI-induced counterfactual wages feed into the economy-wide gender wage gap, we estimate the following regression model:

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

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.

Our simulation framework is deliberately mechanical, taking the 2021 occupational distribution and each occupation’s

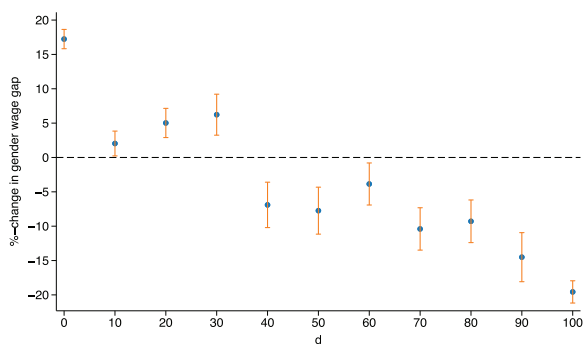
exposure and complementarity to generative AI as given. As a result, we do not model behavioural or general-equilibrium responses to generative AI. The resulting estimates should therefore be interpreted as partial-equilibrium calculations under a given set of assumptions about exposure-aligned wage changes, showing how the aggregate gender wage gap would move if wages changed according to these assumptions while occupational sorting remained fixed.

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.

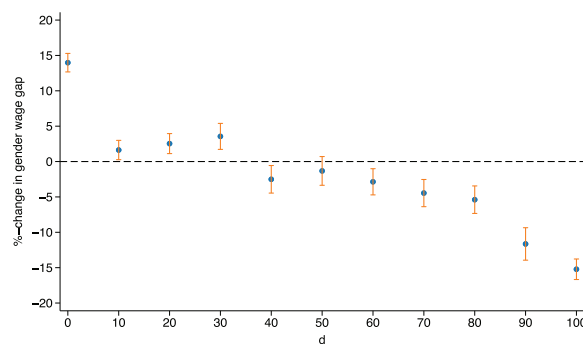
Figure 2a shows results from the baseline linear exposure simulation introduced above. Consider the first estimate at $d = 0$, which assumes that generative AI raises wages in all occupations, and where the wage gain is scaled by the exposure to generative AI (i.e., the individuals with the highest exposure 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 lead to 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 exposure-aligned wage change, the overall effect on the gender wage gap can be considerable, ranging from a decrease of about 17% to an increase of about 20%. 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.

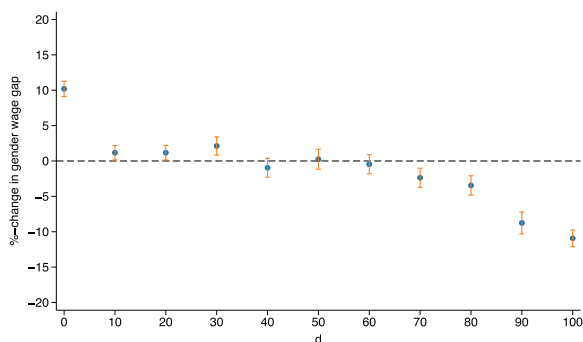
Figure 2a shows the effect on the gender wage gap when keeping the interaction between exposure and wages constant while varying how this interaction changes across the complementarity distribution. As a robustness check, we first run three additional simulations to verify that our findings are not driven by the linearity assumption in Equation (4). These specifications alter the way $\text{Exposure}_{o(i)}$ interacts with the scaling parameter δ . Figure 2b–d present these results. Specifically, Figure 2b replaces $\text{Exposure}_{o(i)}$ with $\text{Exposure}_{o(i)}^2$ in Equation (4), Figure 2c replaces $\text{Exposure}_{o(i)}$ with $\text{Exposure}_{o(i)}^3$, and Figure 2d replaces $\text{Exposure}_{o(i)}$ with $I\text{Exposure}_{o(i)}$, where $I = 1$ at the median of $\text{Complementarity}_{o(i)}$ and $I = 0$ otherwise. We also illustrate what



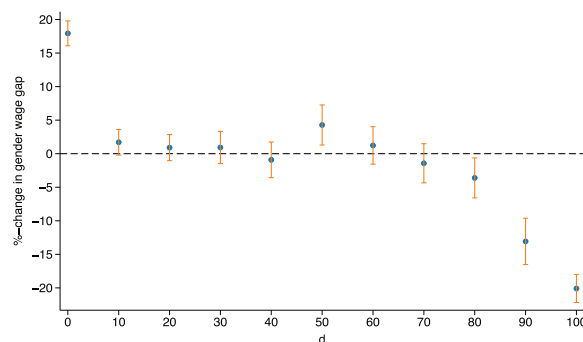
(a) Linear effect across *Exposure*



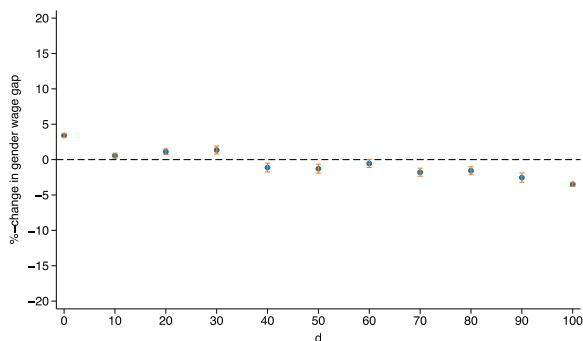
(b) Quadratic effect across *Exposure*



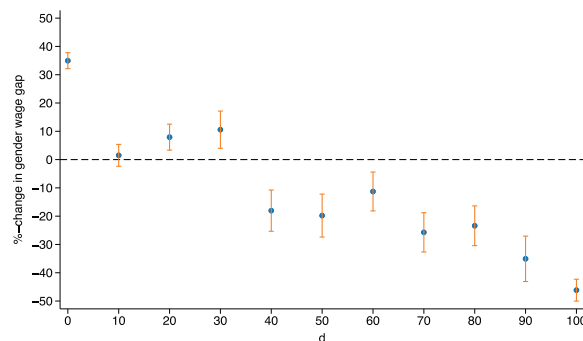
(c) Cubic effect across *Exposure*



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



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



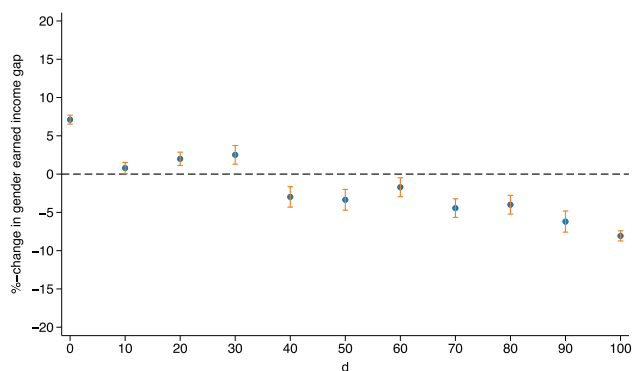
(f) Long-run linear effect across *Exposure*

FIGURE 2 | Simulating changes in the gender wage gap. This figure displays the simulated change in the gender wage gap (y-axis). Positive β_3 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. (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*. [Colour figure can be viewed at [wileyonlinelibrary.com](https://onlinelibrary.wiley.com)]

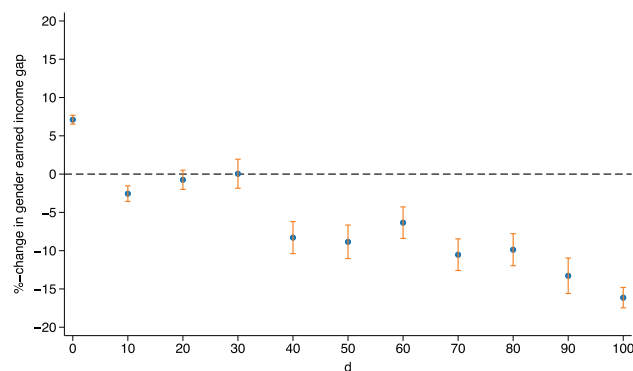
the results look like at different time horizons, where panel (e) displays the association of a 3% annual change after 1 year, and panel (f) the association of a 3% annual change after 10 years, while keeping the exposure term linear. Overall, Figure 2b–f illustrate that the general insights from Figure 2a continue to hold: even under alternative, non-linear exposure mappings or different wage shock magnitudes, the simulated change in the gender wage gap tends to be considerable and, if anything, tends to widen rather than narrow the gender wage gap.

3.4 | Accounting for Heterogeneity by Part-Time Work

Our focus thus far has been on gender differences in full-time-equivalent (FTE) wages, which standardize wages for hours worked. This isolates variation in wage rates but does not directly capture the fact that women are more likely than men to work part-time, which in itself contributes to gender differences in total wage income. Moreover, occupations in



(a) Linear wage adjustment across *Exposure*



(b) Linear wage and hour adjustment across *Exposure*

FIGURE 3 | Simulating changes in the gender earned income gap, accounting for hours worked. This figure displays the simulated change in the gender earned income gap (y-axis), where earnings are defined as the monthly wage scaled by the FTE, with wage and hours adjustment in accordance with Equation (6). Positive β_i 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*(*i*)}, where the wage effect changes from positive to negative, and the hours adjustment changes from null to negative. Panels (a, b) assume a 3% annual wage change over 5 years in the most exposed occupations in our sample in 2021, and panel (b) assumes an additional 3% annual reduction in hours over 5 years for workers in the most exposed non-complementary occupations. (a) Linear wage adjustment across *Exposure*. (b) Linear wage and hour adjustment across *Exposure*. [Colour figure can be viewed at [wileyonlinelibrary.com](https://onlinelibrary.wiley.com)]

which generative AI acts as a substitute may experience not only changes in wage compensation but also reductions in hours worked as parts of the job tasks become automated, further affecting total earnings. To capture these mechanisms, we extend the analysis in two ways.

First, we analyse the simulated earned wage income gap by scaling the FTE wages with workers' contracted hours relative to a 40-h full-time benchmark. This intends to capture gender differences in both wage rates and labour supply. Second, within our simulation framework, we allow hours worked to adjust in occupations that are highly exposed to non-complementary generative AI. Specifically, occupations below the complementarity cutoff d are assumed to not only experience exposure-scaled wage declines but also a proportional reduction in hours worked. In contrast, individuals in high-complementarity occupations receive the same wage-rate premium as in the baseline simulation, but experience no change in hours.⁶

Formally, simulated earned wage income \hat{e}_i is now given by:

$$\hat{e}_i = \begin{cases} w_i(1 - Exposure_{o(i)}\delta) \cdot FTE_i(1 - Exposure_{o(i)}\kappa), & \text{if } Complementarity_{o(i)} \leq d \\ w_i(1 + Exposure_{o(i)}\delta) \cdot FTE_i, & \text{if } Complementarity_{o(i)} > d \end{cases} \quad (6)$$

where w_i is the full-time-equivalent monthly wage, and FTE_i is the individual's observed full-time equivalence (varying from zero to one). The parameter δ governs the wage-rate response to AI exposure, as in the main simulation above, while κ captures the sensitivity of hours worked to exposure in low-complementarity occupations. We assume a cumulated 3% annual reduction in hours over 5 years for occupations that are non-complementary and highly exposed to generative AI. This implies a maximum proportional reduction in hours of about 16% for the most exposed occupations ($\kappa = 1.03^5 \approx 0.159$). This specification is intended as a stylized sensitivity analysis and allows us to assess whether our main conclusions are robust when

accounting also for intensive-margin adjustments in working hours.

Figure 3a shows the baseline linear exposure simulation for the gender earned wage income gap, accounting for part-time work, keeping the labour supply fixed as of 2021, that is, when $\kappa = 0$. The resulting pattern closely resembles Figure 2a, but with somewhat smaller magnitudes. Allowing for adjustments in hours worked in addition to wage adjustments in accordance with Equation (6), as shown in Figure 3b, leads to a further widening of the simulated gender earned income gap relative to the baseline earned income-gap scenario, particularly once the complementarity cutoff exceeds $d = 40$. By construction, no hours adjustment occurs at $d = 0$, as generative AI is assumed to be complementary to the job tasks of all workers. At $d = 50$, workers in the lower half of the complementarity distribution are assumed to experience both exposure-scaled wage losses and proportional reductions in hours worked, while workers in the upper half receive an exposure-scaled wage gain. Under this scenario, the gender earned income gap widens by approximately 9%. Taken together, these results indicate that incorporating hours adjustments reinforces the baseline conclusion that wage and hours reductions in highly exposed, low-complementarity occupations would imply wider gender disparities.

4 | Concluding Remarks

This paper examines how gender-based occupational sorting before 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 alternative assumptions about the relationship between generative AI exposure, task complementarity, and wages across

occupations, holding the 2021 occupational sorting and wage structure fixed. These simulations indicate that exposure-aligned wage changes could have important implications for the gender wage gap and that even modest negative wage changes in a small share of occupations would imply 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 labour 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 assessing possible implications for the gender wage gap.

Acknowledgements

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, Madalena Ronchi, and Björn Tyrefors for excellent discussions and suggestions. We also thank Carlo Pizzinelli, Augustus Panton, Marina Mendes Tavares, Mauro Cazzaniga, and Longji Li for generously sharing the data underlying the AI complementarity index used in this paper. Ludvig Rydén provided excellent research assistance. The replication instructions and code for all figures and tables are available at <https://doi.org/10.1111/obes.25246>. Open access publishing facilitated by Hanken Svenska Handelshögskolan, as part of the Wiley - FinELib agreement.

Funding

This work was supported by Jan Wallander and Tom Hedelius Foundation and the Tore Browaldh Foundation, P25-0107; Marianne and Marcus Wallenberg Foundation, 2024.0019.

Data Availability Statement

The replication instructions and code required to reproduce all figures and tables in this paper are available at <https://doi.org/10.1111/obes.25246>. The underlying Swedish administrative microdata are subject to confidentiality restrictions and cannot be shared publicly. Researchers wishing to replicate our study may apply for data access directly through Statistics Sweden.

Endnotes

¹See for example, Noy and Zhang [1]; Brynjolfsson, Li and Raymond [2]; Eloundou, Manning, Mishkin and Rock [3]; Acemoglu [4].

²Beyond differences in occupational exposure, new evidence shows that women are less likely than men to adopt generative-AI tools [21–23], even though high-performing women may stand to gain the most from using the technology [24]. Consequently, even when generative AI complements tasks and raises wages, gender gaps in adoption patterns could magnify overall wage disparities.

³Note that we use the explicit measure of potential complementarity of AI from Pizzinelli et al. [19], which they label as θ .

⁴In January 2026, the USD to SEK exchange rate was approximately 9.2.

⁵Unreported results show that exposure is higher among more educated workers.

⁶Approximately 76% of employed individuals work full time in our sample; hence, adjusting upwards in hours is less likely. Among part-time workers in 2021, only around 17% report that they would prefer to work more hours [26].

References

1. S. Noy and W. Zhang, “Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence,” *Science* 381, no. 6654 (2023): 187–192.
2. E. Brynjolfsson, D. Li, and L. Raymond, “Generative AI at Work,” (2023), Technical Report, National Bureau of Economic Research, Cambridge, MA.
3. T. Eloundou, S. Manning, P. Mishkin, and D. Rock, “GPTs Are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models,” (2023), *arXiv Preprint arXiv:2303.10130*.
4. D. Acemoglu, “The Simple Macroeconomics of AI,” 2024 NBER Working Papers 32487, National Bureau of Economic Research, Inc.
5. P. Beaudry and E. Lewis, “Do Male-Female Wage Differentials Reflect Differences in the Return to Skill? Cross-City Evidence From 1980–2000,” *American Economic Journal: Applied Economics* 6, no. 2 (2014): 178–194.
6. S. E. Black and A. Spitz-Oener, “Explaining Women’s Success: Technological Change and the Skill Content of Women’s Work,” *Review of Economics and Statistics* 92, no. 1 (2010): 187–194.
7. S. Yamaguchi, “Changes in Returns to Task-Specific Skills and the Gender Wage Gap,” *Journal of Human Resources* 53, no. 1 (2018): 32–70.
8. P. Cortés, Y. Feng, N. Guida-Johnson, and J. Pan, “Automation and Gender: Implications for Occupational Segregation and the Gender Skill Gap,” *Journal of Labor Economics* (2024).
9. C. G. Aksoy, B. Özcan, and J. Philipp, “Robots and the Gender Pay Gap in Europe,” *European Economic Review* 134 (2021): 103693.
10. S. Ge and Y. Zhou, “Robots, Computers, and the Gender Wage Gap,” *Journal of Economic Behavior & Organization* 178 (2020): 194–222.
11. E. W. Felten, M. Raj, and R. Seamans, “Occupational Heterogeneity in Exposure to Generative AI,” *SSRN Electronic Journal* (2023).
12. G. M. Cortes, A. Oliveira, and A. Salomons, “Do Technological Advances Reduce the Gender Wage Gap?,” *Oxford Review of Economic Policy* 36, no. 4 (2020): 903–924.
13. M. Bertrand, “New Perspectives on Gender,” in *Handbook of Labor Economics, Vol. 4 of Handbook of Labor Economics*, ed. O. Ashenfelter and D. Card (Elsevier, 2011), 1543–1590.
14. F. D. Blau and L. M. Kahn, “The Gender Wage Gap: Extent, Trends, and Explanations,” *Journal of Economic Literature* 55, no. 3 (2017): 789–865.
15. C. Goldin, “A Grand Gender Convergence: Its Last Chapter,” *American Economic Review* 104, no. 4 (2014): 1091–1119.
16. C. Olivetti and B. Petrongolo, “The Economic Consequences of Family Policies: Lessons From a Century of Legislation in High-Income Countries,” *Journal of Economic Perspectives* 31, no. 1 (2017): 205–230.
17. I. Sin, S. Stillman, and R. Fabling, “What Drives the Gender Wage Gap? Examining the Roles of Sorting, Productivity Differences, Bargaining, and Discrimination,” *Review of Economics and Statistics* 104, no. 4 (2022): 636–651.
18. P. Gmyrek, J. Berg, and D. Bescond, “Generative AI and Jobs: A Global Analysis of Potential Effects on Job Quantity and Quality,” 2023 ILO Working Paper, 96.
19. C. Pizzinelli, A. J. Panton, M. M. Tavares, M. Cazzaniga, and L. Li, “Labor Market Exposure to AI: Cross-Country Differences and Distributional Implications,” 2023 IMF Working Papers, (23/216).
20. M. Cazzaniga, A. Panton, L. Li, C. Pizzinelli, and M. M. Tavares, “A Gender Lens on Labor Market Exposure to AI,” in *AEA Papers and Proceedings*, vol. 115 (American Economic Association, 2025), 56–61.
21. D. Carvajal, C. Franco, and S. Isaksson, *Will Artificial Intelligence Get in the Way of Achieving Gender Equality?* Discussion Paper Series in Economics No. 3/2024 (2024).

22. A. Humlum and E. Vestergaard, "The Adoption of ChatGPT," *SSRN Electronic Journal* (2024).
23. N. G. Otis, K. Cranney, S. Delecourt, and R. Koning, "Global Evidence on Gender Gaps and Generative AI," 2024 OSF Preprints h6a7c, Center for Open Science.
24. C. Franco, N. Irmert, and S. Isaksson, *Will AI Help or Hurt Learning?* (Mimeo, 2025).
25. E. Felten, M. Raj, and R. Seamans, "Occupational, Industry, and Geographic Exposure to Artificial Intelligence: A Novel Dataset and Its Potential Uses," *Strategic Management Journal* 42, no. 12 (2021): 2195–2217.
26. SCB, "Arbetsutbud Och Jobbyten Under Inflationsperioden 2022–2023," 2025 Rapport 2025:3, Statistiska centralbyrån, Sweden.