The potential impact of AI on the gender gap.

Arnaud Dupuy * Camila Pineda León *

March 20, 2025

Abstract

This paper analyzes the impact of Artificial Intelligence (AI) on the gender wage gap through its effects on estimated values of job amenities and perceived productivity. We employ a one-to-one matching model using data from the UK for 2021-2022 and an AI occupational exposure score. Our results indicate that the gender gap in perceived productivity is nearly 4% in occupations highly exposed to AI but expands to 18.21% in low AI-exposed jobs. Furthermore, job amenities values for workers decrease by 0.11 log-points in intensely AI-focused occupations, with this value being 2.2% lower for female workers.

*University of Luxembourg

1 Introduction

The rapid acceleration of Artificial Intelligence (AI) and its effects and applications are sparking excitement and concern in societies. AI was born in 1950 and its journey has been characterised by breakthroughs and leaps. In the 21st century, the explosion of data science and computational advancements have enabled AI to tackle a wider range of problems. By 2010, deep learning algorithms had achieved state-of-the-art performance in tasks such as image recognition and natural language processing. In some instances, deep learning neural networks have been shown to perform better than humans at classifying images, as evidenced by lower error rates in benchmark tests (Dodge and Karam (2017)). The early 2020s have witness the emergence of Generative Pre-trained Transformers (GPTs) in the market which are being extensively employed by many users. For example, ChatGPT, the GPT model released by OpenAI in November 2022, had already 100 millions of active users in January 2023.

The capacity of AI to perform a broad spectrum of tasks generate fear among the workforce as seen in previous automation waves. Interestingly, one characteristic of AI different to other technologies is its ability to undertake tasks traditionally considered as cognitive tasks. One known example of this are radiologist, whose working activities seem to be highly exposed to AI (Hosny et al. (2018)). This points to the fact that AI applications surpass the automation of mere routine tasks, which predominantly characterize lower-wage occupations.

Despite the potential importance of AI and its effects on the labor market, research into the impact of AI on occupations by demographic group remains limited. Our focus is particularly on gender differences, as gender gaps persist. In the UK, the gender pay gap, measured by median gross hourly earnings, was 14.9% in 2022. While there is extensive research on the gender pay gap and the effects of AI on employment and wages, studies focusing on the implications of AI for the gender gap are scarce. In this paper, we aim to address this gap in the literature with a novel approach. We employ a static structural model that provides a bipartite view of the labor market. This enables us to estimate the value of two non-observed variables in the data: job amenities and perceived productivity. In our model, these two variables affect equilibrium wages and matching probabilities. Therefore, the impact of AI on the gender gap can be analyzed through the effects of AI on workers' job amenities and firms' perceived productivity. Specifically, female and male workers might value occupations more exposed to AI differently, while firms might perceive varying levels of productivity from male and female workers in highly AI-intensive jobs. Our paper structure is as follows: Section II explores the literature related to our topic, Section III describes our sample; Section IV presents

the methodology; Section V discusses the results; Section VI present some counterfactual estimations and Section VII concludes.

2 Literature review

Previous literature has analysed technological effects on labor. D. Autor and Dorn (2013) and Brussevich, Dabla-Norris, and Khalid (2019) found that automation has led to occupational polarisation in the US, characterized by a shift away from middle-skilled roles towards higher- and lower-skilled positions. This redistribution appears to have relatively benefited women in terms of job losses, as the female labor force was predominantly concentrated in low-paid and low-skilled occupations. Aksoy, Özcan, and Philipp (2021) provided evidence that robotization has led to an increase of gender gap in Europe, as middle-skilled and high-skilled male workers have increased their productivity compared to female workers, especially in countries with high initial levels of inequality. Evaluating the effects of AI on the labor market requires a measure of AI's penetration within occupation. This task is challenging as there is still not of a consensus on AI's definition and scope.

Some authors have developed some metrics to quantify the AI exposure or AI content in occupations (Brynjolfsson, Mitchell, and Rock (2018), Felten, Raj, and Seamans (2019), Felten, Raj, and Seamans (2021), Felten, Raj, and Seamans (2023), Fernández-Macías and Bisello (2017), Martínez-Plumed et al. (2020), Tolan et al. (2021), Webb (2019), Zarifhonarvar (2023)). Among them, some researches (Brynjolfsson, Mitchell, and Rock (2018), Tolan et al. (2021) and Webb (2019)) have identified a positive correlation between the level of AI exposure in occupations and wages. This relationship is attributed to AI's capacity to perform cognitive tasks, which include pattern detection, optimization or analysis. High-skilled occupations, which require relatively higher volume of cognitive tasks compared to middle- and low-skilled occupations that involve more routine tasks, are the most affected. Webb (2019) develops several indicators for occupational exposure to robots, software and AI technologies. While robotization predominantly affects occupations in the lowest wage percentile, software impacts middle-paid occupations, and AI more significantly affects those in higher wage percentiles. This evidence suggests that AI impacts jobs that were not directly affected in previous waves of automation.

Using an indicator developed by Felten, Raj, and Seamans (2021), the research of Albanesi et al. (2023) and Guarascio, Reljic, and Stöllinger (2023) studied the effects of AI on employment and concluded that AI has positive effects on European regional employment. Stapleton,

Copestake, and Pople (2021) analyzed 15 million job vacancies posts in India and found that since 2016, the increase in AI skills demands from firms has reduced the growth in non-AI job posts and the average wage offered in the lowest paid roles. Additionally, Webb (2019) presents a positive correlation between the proportion of female workers in an occupation and its AI exposure score in the US, highlighting the disparate effects of AI across different segments of the job market.

Other studies have focused on the gap in AI adoption, specifically using the case of ChatGPT. In an experiment with students from Norway, Carvajal, Franco, and Isaksson (2024) found that male students achieved a higher success rate than female students in writing prompts -34%- and estimated a gender gap in ChatGPT adoption of 15 percentage points. These results are driven by a higher perception among female students of using AI tools in learning as cheating, and by greater persistence of male students when they do not receive the desired answer from ChatGPT. Moreover, there is a premium on entry-level jobs for top-graded female students (but not for male students) who posses AI skills, compared to female students without such skills. This highlights the positive effects of AI knowledge on hiring opportunities for women. Humlum and Vestergaard (2024) conducted an extensive survey in Denmark and calculated a 20 percentage point gender gap in ChatGPT adoption among workers in occupations exposed to AI, attributed to a higher likelihood of women facing adoption barriers.

Another important branch of literature for our research is related to the reasons behind the gender pay gap. The increasing participation of women in the labor market and the attainment of equal or greater years of education by women compared to men in developed countries have diminished the impact of human capital differences on the gender pay gap. However, gender differences in pay persist, and there is extensive literature exploring the reasons behind this. Blau and Kahn (2017) highlight the interruptions in women's careers as one of the factors that continue to contribute to the pay gap, especially in high-paying occupations. Regarding this, Goldin (2014) remarks the importance of working continuously long hours in some occupations where flexibility is viewed as a cost. Consequently, mothers who require greater flexibility to balance their family and work lives face greater challenges in occupations that highly reward long and specific working hours.

Other significant reasons contributing to the gender pay gap include occupational and industrial factors, as highlighted by Blau and Kahn (2017) and Blau and Kahn (2020). Industry and occupation accounted for the 51% of the gender gap in the US in 2010, making them the

largest measurable variables to explain the gap. Recent research has explored the effects of personal traits on the gender gap. For instance, Mueller and Plug (2006) found that differences in the agreeableness trait between men and women are the most significant trait in explaining gender differences in earnings in the US. However, the impact of personality trait differences is a residual factor compared to child penalties and occupational and industry explanations behind gender gap (Blau and Kahn (2020)).

3 Data

3.1 Understanding Society survey

The study is conducted in the United Kingdom. We employ the survey "Understanding Society". This annual survey interviews to 40,0000 households, gathering data on demographics, family dynamics, employment, and personality traits of the participants. We choose for our analysis the last available wave of data, corresponding to the 2021-2022 period. Our sample will be restricted to employees aged between 25 and 65 years old and who work at least 30 hours per week. This resulted in a sample size of 6,816 observations.

Table 1: Descriptive statistics of workers characteristics.

	All	Women	Men
% Women	47.73		
Average age	45.02	44.63	45.38
Average yrs of experience	25.36	24.80	25.87
% Higher education	60.75	65.02	56.86
% Single	16.65	19.06	14.45
% Parent	35.02	30.99	38.70
% White	83.79	82.45	85.01
Average log hourly wage	2.81	2.74	2.87

Note: Years of experience refers to years of potential experience computed as age minus years of education minus five. Higher education denotes if the individual has completed bachelor degree or higher. Parent is defined as either the biological or adoptive mother or father of a child under 16 years old.

Table 1 presents the average values of the main worker characteristics in our sample. Women constitute nearly 48% of the observations. On average, out of every 100 women in our sample, 65 have completed Bachelor's studies or a higher level of education. For men, the average is 57 out of 100 who have graduated with a Bachelor's degree or higher. This represents a gender gap in higher educational attainment of -14.35%, aligning with trend of the gain in human capital among female workers compared to male workers. In terms of potential years of experience, men in our sample have an average of 25.87 years of experience,

which is one year more than their female counterparts. Male workers earn, on average, 4.5% higher wages than women as measured by log-hourly wages.

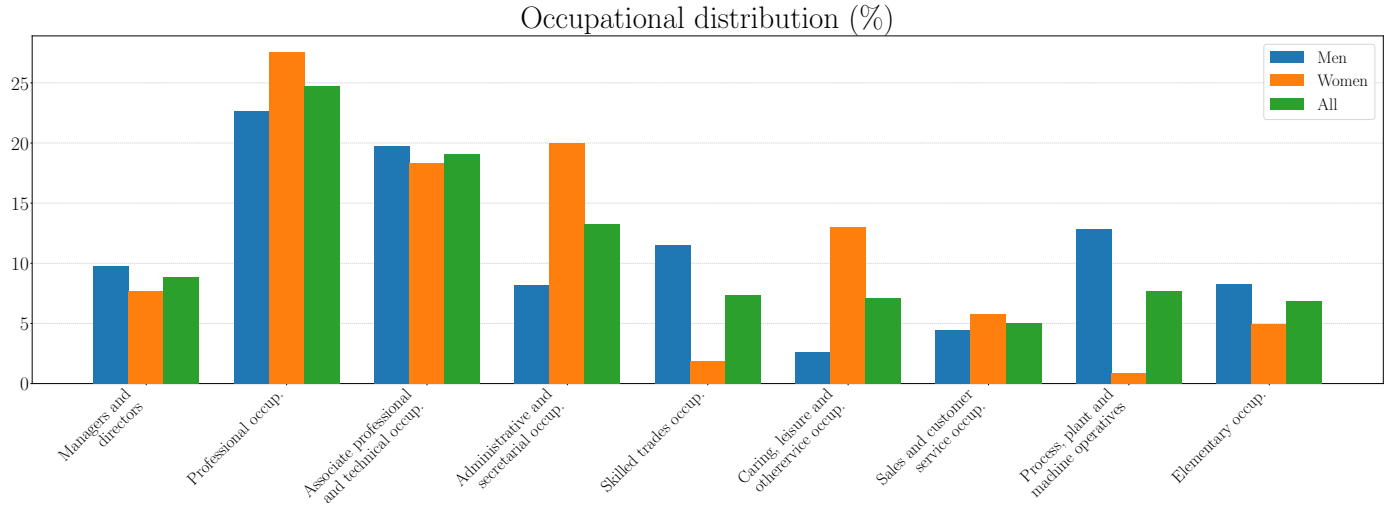


Figure 1: Occupational distribution

In Figure 1, we plotted the occupational distribution of female, male, and all workers. Notice that the distribution across occupations for employed women and men is heterogeneous. Professional occupations employ 20% of our sample, with this percentage being higher for female workers than for male workers. The graph also highlights different gender sorting patterns into occupations. For instance, 'Skilled trade occupations' and 'Process, plant and machine occupations' are predominantly male dominated, while 'Caring, leisure and other service occupations' and 'Administrative and secretarial occupations' have a larger proportion of female workers.

3.2 Artificial Intelligence Occupational Exposure score

To assess the impact of AI across occupations, we rely on an exposure score developed by Felten, Raj, and Seamans (2021). This Artificial Intelligence Occupational Exposure (AIOE) score measures the extent to which an occupation's abilities are exposed to AI applications. To construct this AI measure, the authors collected information on AI applications, working abilities, and occupations. Felten, Raj, and Seamans (2021) utilized the AI Progress Measurement project from the Electronic Frontier Foundation (EFF) to select a set of ten AI applications: abstract strategy games, real-time video games, image recognition, visual question answering, image generation, reading comprehension, language modeling, translation, speech recognition, and instrumental track recognition. Additionally, they obtained information on abilities from

the O*NET database, a resource developed by the United States that provides data on the importance and prevalence of 52 abilities used in occupational activities.

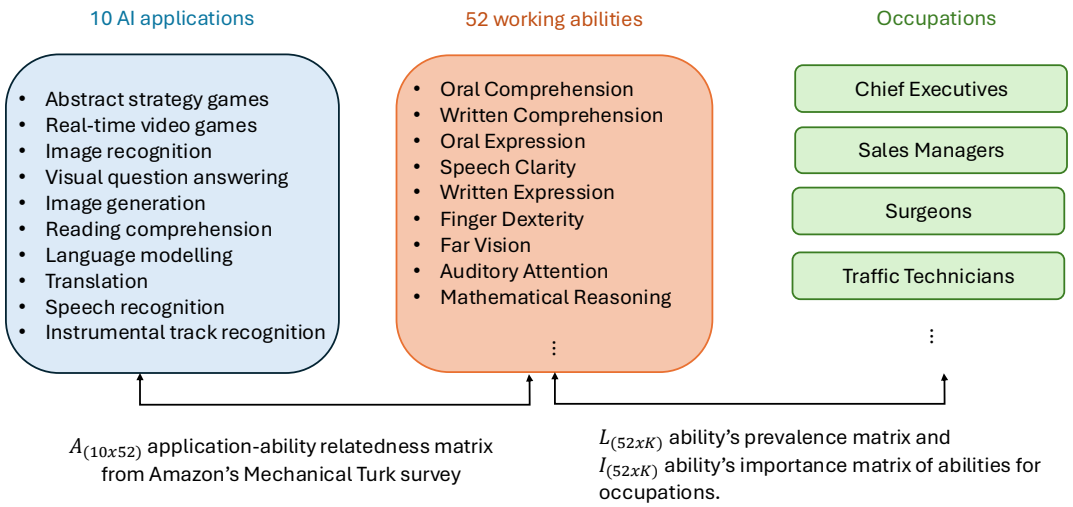


Figure 2: AI applications, abilities and occupations mapping.

They link AI applications to abilities using an application-ability relatedness matrix, denoted as the A matrix, obtained from Amazon's Mechanical Turk survey. The cells of the A matrix, denoted as $x_{i,j}$, have values between 0 and 1 that reflect ratings on how much the j-th ability is exposed to the i-th AI application. The total sum over rows, $\sum_{i=1}^{10} x_{ij}$, gives the exposure value of the j-th ability to AI. The next step is to map these ability exposure values to occupations. We denote L and I as the matrices that collect the level of prevalence and the level of importance of these 52 abilities for occupations. We mapped the SOC occupational code from the US, as used in the O*NET dataset, to the SOC nomenclature used in the 'Understanding Society' survey to obtain a standardized AIOE measure for the UK. As the correspondence between the O*NET dataset and the 'Understanding Society' survey occupations is not one-to-one, we compute the average values of the level of prevalence and the level of importance for each ability that links to each occupation in the SOC UK nomenclature. This results in two matrices L and I of size (52,104) that contain a single value of the level of prevalence and the level of importance of the j-th ability for the k-th occupation. We compute the AIOE of an occupation following Felten, Raj, and Seamans (2021) as:

$$AIOE_k = \frac{\sum_{j=1}^{52} AI_{ij} \times L_{jk} \times I_{jk}}{\sum_{j=1}^{52} L_{jk} \times I_{jk}}$$

Our final variable AIOE contains the standarised values of all the $AIOE_k$ of our dataset. It is important to note that with this AI indicator, we cannot account for complementary or substitutive effects of AI on occupations, but only how intensively an occupation is exposed to AI. Figure 2 illustrates this mapping process between AI applications, working abilities, and occupations.

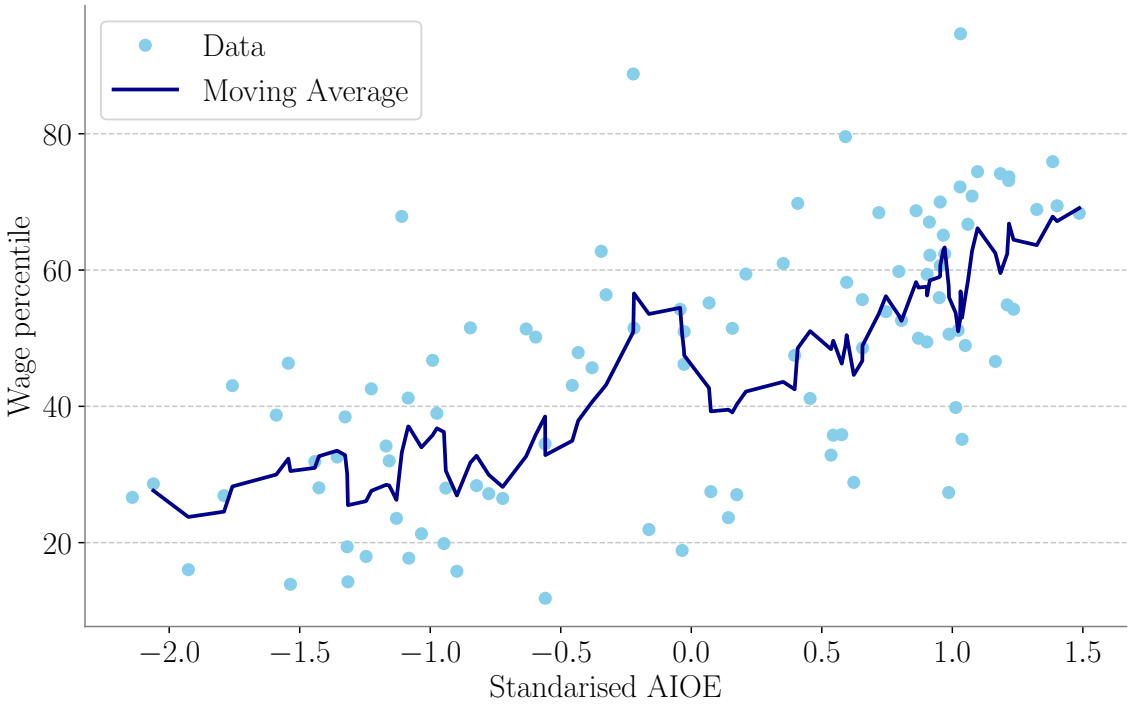


Figure 3: AIOE by wage percentile

In Tables 2 and 3 we provide examples of abilities' exposure to AI and the importance of these abilities for two selected occupations. Table 2 displays two abilities along with their descriptions and AI exposure values. The AI exposure values for the *information ordering* and *arm-hand steadiness* abilities are 8.775 and 2.672, respectively. While *information ordering*—a cognitive ability—is 3.3 times more exposed to AI than *arm-hand steadiness*, a psycho-motor ability, this example highlights that cognitive abilities tend to be more affected by AI than other types of working abilities.

In Table 3, we present the importance scores (on a scale from zero to five) for these selected abilities across two occupations. For *computer network architects*, the ability *information ordering* is 1.5 times more important than it is for *painting, coating, and decorating workers*.

Conversely, *arm-hand steadiness* is 1.8 times more relevant for *painting, coating, and decorating workers* than for *computer network architects*. Therefore, when focusing solely on these two abilities and occupations, since *information ordering* is both more exposed to AI and more crucial for *computer network architects* than for *painting, coating, and decorating workers*, the former occupation is more exposed to AI than the latter.

Table 2: Sum of AI applications' scores by ability

Ability	**Description**	$\sum_{i=1}^{10} x_{ij}$
Information ordering	Arrange things or actions in a certain order or pattern according to a specific rule or set of rules	8.775
Arm-Hand Steadiness	The ability to keep your hand and arm steady while moving your arm or while holding your arm and hand in one position	2.672

Table 3: Importance of abilities for occupations

Occupation	**Ability**	Importance level (0-5 scale)
Computer Network Architects	Information ordering	4
	Arm-Hand Steadiness	2
Painting, Coating, and Decorating Workers	Information ordering	2.63
	Arm-Hand Steadiness	3.63

In table 4 we show the ten occupations with the highest and lowest AIOE in the UK. Notice that as according to other studies, the most impacted occupations are white collar occupations, contrary to what we have witness in previous automation waves. Furthermore, Figure 3.2 suggests a positive correlation between wage percentiles and the AIOE metric across occupations.

Table 4: Occupations with the highest and lowest values of AIOE.

Highest AIOE	**Lowest AIOE**
Chartered and certified accountants	Sports players
Actuaries, economists and statisticians	Roofers, roof tilers and slaters
Psychologists	Painters and decorators
Barristers and judges	Plasterers
Taxation experts	Vehicle valeters and cleaners

4 Methodology

To explore the potential impact of Artificial Intelligence (AI) on gender gap within the UK labor market in 2021-2022, we adopt a bipartite matching model with logit unobserved heterogeneity as proposed by Dupuy and Galichon (2022). This framework considers the two sides of the labor market in equilibrium (workers and firms) and models how workers match with their optimal firm and how firms hire the optimal worker given wages. We assume that both agents act as price-takers and utility maximizers. The advantage of this approach is the possibility of estimating the values of two non-observed parameters: job amenities and perceived productivity [1].

There are a list of variables that we observe in our data. The set of characteristics of workers we examine, are represented by a vector of attributes $x \in X = \mathbb{R}^{d_x}$ where d_x denotes the dimensionality of the worker attribute space. Similarly, characteristics of firms are contained in a vector of attributes $y \in Y = \mathbb{R}^{d_y}$ with d_y representing the dimensionality of the firm attribute space. All workers i and firms j belong to a type x of worker and a type y of firm, respectively. Additionally, our dataset also include information of the transfers in equilibrium (wages) and probability density of occurrence of a match between a worker of type x and a firm of type y. We denote with $w(x, y)$ to the wage perceived by a worker of type x working for a firm type y and $\pi(x, y)$ to the feasible matching between worker x and firm y. There are other variables non-observed in the data but that our agents value into their utility maximizing problems. Worker of type x values not only the wage but also the job amenities they perceive when working for a firm y. Firm y considers the wage and the perceived productivity she has for that type of worker x during her hiring process.

Workers' utility maximizing problem.

Each worker i of this market has a type x and aims at maximising her utility when working for a firm of type y. When maximizing her utility function a worker i takes into account the following two components: i) wages and, ii) job amenities. This optimisation problem is modeled as:

$$\max_{y \in Y} \{\alpha(x, y) + w(x, y) + \sigma_1 \epsilon(y)\}$$

Wages $w(x, y)$ are the monetary transfer between a pair (x, y) in equilibrium and job amenities are captured by the sum $\alpha(x, y) + \epsilon(y)$. Firm y characteristics that are valued by worker

[1] Perceived productivity refers to true productivity of workers and employers' taste discrimination productivity following to Becker (2010).

x are included in $\alpha(x, y)$. Notice that there is a deterministic part in this utility function, $w(x, y) + \alpha(x, y)$, that is the same for all workers of type x working for firms of type y. The stochastic part $\epsilon(y)$ captures the specific preference for the amenities of firm y that a particular worker has for a job when maximising her utility. This worker-specific random process $\epsilon(y)$ is modelised as a Gumbel random process and is scaled by a factor σ_1. Notice that in the optimization problem, each worker chooses between types of firm from Y and therefore, she will be indifferent between firms that belong to the same type $y \in Y$.

Firms' profit maximizing problem.

Each firm j of this market has a type y and incur in a cost given by $w(x, y)$ for hiring a worker of type x. Additionally, firms also value the perceived productivity of their employees as collected by the sum $\gamma(x, y) + \eta(y)$. Symmetrically to workers, the deterministic part of the profits for a firm of type y when hires a worker of type x is given by $\gamma(x, y) - w(x, y)$ while the Gumbel-random process $\eta(y)$ represents the particular firm' perceived productivity when matching a worker of type x. The firm' profit maximisation problem is as follows:

$$\max_{x \in X} \{\gamma(x, y) - w(x, y) + \sigma_2 \eta(x)\}$$

where σ_2 is a scaling factor. As before, the firm selects the type of worker x between all the type of workers contained in X and will not have a preference between workers of the same type.

Equilibrium.

The equilibrium of this model is reached when the wage curve is such that the feasible matching from worker's and firm's problem are the same. The equilibrium in this model is defined by the following expressions for wage and feasible matching:

$$\pi(x, y) = \exp\left(\frac{\alpha(x, y) + \gamma(x, y) - u(x) + \sigma_1 \log\, f(x) - v(y) + \sigma_2 \log\, g(y)}{\sigma}\right) \tag{1}$$

$$w(x, y) = \frac{\sigma_1}{\sigma}(\gamma(x, y) - v(y) + \sigma_2 \log\, g(y)) + \frac{\sigma_2}{\sigma}(u(x) - \sigma_1 \log\, f(x) - \alpha(x, y)) \tag{2}$$

where σ_1, σ_2 are non-negative and $\sigma = \sigma_1 + \sigma_2$. We define $f(x)$ as the probability density of a worker to belong to a particular type characterized by the attribute vector x and $g(y)$ as the probability density of a firm to belong to a particular type characterized by the attribute vector y. $u(x)$ and $v(y)$ represent the indirect utility of a worker of type x and the indirect

utility of a firm of type y, respectively. For simplicity and as the entire explanation of the model can be found in Dupuy and Galichon (2022), we have skipped the algebra to reach the equilibrium from the utility maximizing problems. We presented equilibrium equations to be able to focus on the identification of our non-observed values of job amenities and perceived productivity.

Note that in equations 1 and 2 we express two observed variables in the data, $w(x, y)$ and $\pi(x, y)$ as a function of non-observed $\alpha(x, y)$ and $\gamma(x, y)$ variables. While $\gamma(x, y)$ is associated with a positive sign in equations 1 and 2, $\alpha(x, y)$ contributes positively to equation 1 but negatively to equation 2. This differing impact on the feasible matching and wage equations enables the identification of $\alpha(x, y)$ and $\gamma(x, y)$. These different effects are reasonable as a worker of type x who is perceived by firms of type y as more productive compared to other types of workers is more likely to be offered a higher wage and to be hired with a greater probability. Hence, both the wage and the probability of the match between a pair of worker and firm (x, y) increase with the perceived productivity. Conversely, if a worker of type x values the job amenities at a firm of type y more than those offered by other firms, they may be willing to accept a lower wage, as their valuation is not solely based on the wage. At the same time, if worker of type x highly values the amenities offered by firm type y above all others, she is more likely to match with a firm of type y.

Parametrization.

The values of $\alpha(x, y)$ and $\gamma(x, y)$ are linearly parametrized and estimated using maximum log-likelihood, as explain in the next section on estimation. Our preferred specification for the systematic value of job amenities and perceived productivity is as follows:

$$\alpha(x, y, A) = \tau(\text{AIOE} \times \text{women}) + \lambda Y \tag{3}$$

$$\gamma(x, y, \Gamma) = \rho(\text{AIOE} \times \text{women}) + \kappa X \tag{4}$$

where $A = (\tau, \lambda)$ and $\Gamma = (\rho, \kappa)$.

The X matrix contains observed attributes of workers in the data. This includes dummy variables for gender (women), higher education, race (white), marital status (single), parenthood, and a series of nine dummy variables for experience intervals spanning five years, from 0 to 50 years of experience, with the reference category being 0-5 years of experience. The Y matrix contains attributes of firms. It includes the AIOE score, which is a continuous variable,

a dummy variable for private firms, occupational dummies, and a 'size' dummy variable that takes the value one if the firm employs more than 200 workers.

Therefore, when valuing the amenities α of a job, workers consider all the characteristics of firms contained in Y, plus an interaction term between AIOE and women that captures whether working in highly AI exposed jobs changes women's job amenities values compared to men's workers. The perceived productivity of firms, γ, is explained by the workers' characteristics in the X matrix, plus the term AIOE×women, which captures a potentially different perceived productivity for female workers in AI-intensive occupations. λ and κ are vectors of coefficients representing the contributions of firm attributes to job amenity values and worker attributes to perceived productivity, respectively. Our coefficients of interest, τ and ρ, reflect how female workers value jobs with high AI intensity differently than male workers and how firms in occupations with high AI exposure perceive the productivity of female workers compared to male workers in the same occupations.

Since our wage and matching equations depend on the α and γ values, the τ and ρ coefficients contribute to explain wage levels and wage differentials in our market equilibrium. This enables us to analyze the effects of our AIOE variable and its interaction with women on wages, through our estimated values of job amenities and perceived productivity.

Estimation.

To fully estimate the model is necessary to estimate the value of the parameters $A, \Gamma, \sigma_1, \sigma_2, t, s^2$. We estimate the vector of parameters $\theta = (A, \Gamma, \sigma_1, \sigma_2, t, s^2)$ using maximum likelihood approach until the predicted moments of the matching distribution and predicted wages equal their sample counterparts. The steps to achieve this are:

1. Initialize the problem choosing random values for Φ, t, σ_1 and σ_2.
2. Calculate ϕ.
3. Compute the predicted matching values as described by equation 5.
4. Compute the predicted wages as described by equation 6.
5. Calculate log-likelihood function value $\log \hat{L}(\theta)$ and their correspondient gradients.
6. Maximise $\log \hat{L}(\theta)$ using Quasi-Newton method.

The first step is straightforward. Assigning random values to vector $\Phi = (A, \Gamma)$ allow us to compute initial values of job amenities and perceived productivity as defined by equations 3

and 4. The second step consists in computing the surplus of the matches (x, y). This surplus is denoted by ϕ and is the result of the sum $\alpha(x, y, A) + \Gamma(x, y, A)$. The third and fourth steps consist in computing the sample estimates of equations 1 and 2 and that are defined as:

$$\pi_{ij}(\Phi) = \exp(\phi_{ij}(\Phi) - a_i(\Phi) - b_j(\Phi)) \tag{5}$$

$$w_i(\theta) = \sigma_1(\gamma_{ii}(\Gamma) - b_i(\Phi)) + \sigma_2(a_i(\Phi) - \alpha_{ii}(A)) + t \tag{6}$$

where in this sample analog a_i and b_j are obtained as the solution to this system of equations:

$$\begin{cases} \sum_{j=1}^{n} \exp\left(\phi_{ij}(\Phi) - a_i - b_j\right) = \frac{1}{n}, & \forall i = 1, \ldots, n, \\ \sum_{i=1}^{n} \exp\left(\phi_{ij}(\Phi) - a_i - b_j\right) = \frac{1}{n}, & \forall j = 1, \ldots, n, \end{cases} \tag{7}$$

In step fifth, we compute $\log \hat{L}_1(\theta)$ and $\log \hat{L}_2(\theta)$ as the log-likelihood functions of equations 6 and 5. These are defined as:

$$\log \hat{L}_1(\theta) = \sum_{i=1}^{n} (\phi_{ii}(\Phi) - a_i - b_j) \tag{8}$$

and

$$\log \hat{L}_2(\theta) = -\sum_{i=1}^{n} \frac{(W_i - w_i(\theta))^2}{2s^2} - \frac{n}{2} \log s^2 \tag{9}$$

Therefore, the goal is to maximise the log-likelihood of our sample given by the function:

$$\log \hat{L}(\theta) = \log \hat{L}_1(\theta) + \log \hat{L}_2(\theta) \tag{10}$$

where s^2 is the mean squared error from equilbrium wage equation. From the log-likelihood functions we can compute the partial derivatives with respect to all the parameters in θ as explained in Appendix A. Using both, log-likelihood function value and its gradients, BFGS method updates the values of vector θ and the process starts again and is repeated until convergence of the predicted moments with their sample counterparts.

Table 5: Effects of firms' and workers' characteristics on job amenities and perceived productivity (2021-2022).

	AIOE×women	Higher educ	Women	White	Single	Parent
Perceived productivity	0.1429	0.2317	-0.1821	0.0062	-0.0434	0.0342
	(0.0146)	(0.0108)	(0.0118)	(0.0122)	(0.0105)	(0.0309)
	AIOE×women	Private	AIOE	Size		
Job amenities	-0.0223	-0.0909	-0.1089	-0.1722		
	(0.0113)	(0.0098)	(0.0112)	(0.0106)		
Salary constant	2.4388					
	(0.0100)					
Occupation controls	Yes					
Experience controls	Yes					
Sigma 1	0.0010					
Sigma 2	0.3341					
R-squared	0.3855					
N	6,816					

Note: All effects are measured in log wages. Standard errors, in parenthesis, are calculated from the Hessian of the likelihood. Salary constant is measure as log hourly wage.

5 Results

Table 5 presents the outcomes from parameterizing $\alpha(x, y, A)$ and $\gamma(x, y, \Gamma)$. The table should be read as follows: the first row contains the coefficients ρ and κ of workers' attributes that parametrize perceived productivity and the second row contains the coefficients τ and λ of firms' attributes that parametrize job amenities value.

From the perspective of perceived productivity of the firms, the coefficient related to AIOE×women shows that firms with occupations with one standard deviation units higher of AIOE measure, perceive women as more productive their male counterpart by 14.19%, being this coefficient statistically significant. We can refer to this positive coefficient as the'AI premium' perceived by female workers. However, this does not offset the negative perceived productivity of female workers compared to male workers, which is lower by 18.21%. It is important to clarify that we refer to perceived productivity, as defined by our model's specifications and that our coefficients should be interpreted with care as they might reveal discrimination. Higher education is associated with an increase in perceived productivity by 0.2317 log-points, being this value statistically significant. Single employees are perceived as less productive by 0.0427 log-points than non-single workers, and white workers are perceived as more productive than non-white workers but this difference is not significant. Interestingly, parents are not perceived as less productive by firms, as reflected by the positive, though not

statistically significant, coefficient of that variable.

In relation to the results for the systematic part of job amenities, we observe that the coefficient of the AIOE variable is negative and significant. This indicates that an increase in AIOE by one standard deviation unit reduces the job amenities value by -0.1089 log-points. Interestingly, the analysis suggests that women perceive a lower amenity value when working for firms with higher AIOE, as indicated by the coefficient for AIOE×women (-0.0223). The coefficients associated to Private and Size variables are negative and statistically significant, implying that public sector jobs and firms with 200 or fewer employees increase the job amenities values of workers compared to private firms and firms with more than more than 200 employees.

The channels through which AI affects job amenities and perceived productivity remain to be explored. One potential factor for the negative impact of AIOE on job amenities values could be that high AI intensity occupations require additional learning, which might negatively affect workers' preferences. Female workers, whose amenities value decreases by an additional 2.2% in highly AI-exposed jobs compared to men, might face higher adoption barriers, in line with the findings of Humlum and Vestergaard (2024). The positive effect of AIOE on the perceived productivity of firms for female workers in AI-intensive occupations could be due to women's quicker adaptation to AI technologies (once they have acces to them), or it may be related to the premium on AI knowledge that women possess, as indicated by Carvajal, Franco, and Isaksson (2024).

6 Counterfactuals

Table 6: Summary statistics of observed and simulated values

	Observed	$\tau = 0$	$\rho = 0$	$\kappa_{\text{women}} = 0$	AIOE
Average wage	17.2757	16.9661	17.5574	15.5343	17.0738
Average women's wage	16.0432	15.8024	15.9627	15.9630	15.8661
Average men's wage	18.4010	18.0286	19.0133	15.1430	18.1764
% Gender Gap	12.8137	12.3483	16.0448	-5.4151	12.7107
Perc. Change av. wage		-1.7922	1.6304	-10.0798	-1.1689
Perc. Change women's av. wage		-1.5009	-0.5012	-0.5000	-1.1039
Perc. Change men's av. wage		-2.0238	3.3275	-17.7056	-1.2209
Perc. Change Gender Gap		-3.6323	25.2158	-142.2605	-0.8039

Note: Average wages as measured in pounds per hour. Counterfactuals values are presented in columns $\tau = 0$, $\rho = 0$, $\kappa_{\text{women}} = 0$, and AIOE .

In this section, we present the results of several counterfactual simulations. We consider

four simulated scenarios, and Table 6 summarizes the statistics for each. The first column displays the observed average hourly wages and the corresponding gender gap, while columns 2 through 5 show the simulated outcomes.

For our first counterfactual, we simulate wages under the assumption that there is no gender difference in preferences for AI job intensity. Using the parameter estimates from Table 5 but setting τ to zero (instead of -0.0223), we compute equations 5 and 6. The results indicate that if both men and women equally value the benefits derived from AI exposure at work, the average hourly wage would decrease by 1.8%. This decrease is more pronounced for male workers (-2%) than for female workers (-1.5%), reducing the gender wage gap from 12.8% to 12.4%. The fall in mean wages is expected since by setting $\tau = 0$ instead of $\tau = -0.0223$, the negative pressure on wages from female workers who perceive a disutility from working in highly AI-exposed jobs disappears. As mentioned, male wages are more affected in relative terms, but they still have a higher average hourly wage (£18.03) than female wages (£15.80), causing the gender gap to continue falling while remaining positive.

For our second counterfactual, we retain the estimated values of θ from Table 5 but change the coefficient of ρ from 0.1429 to 0. This adjustment implies that, regardless of the degree of AI exposure, firms would perceive women to be 18.21% less productive than men. This female penalty in perceived productivity causes the average female wages to fall by 0.5% while male wages increase by 3.3% on average. Although the drop in female wages is less severe compared to the first counterfactual (with $\tau = 0$), the substantial increase in male wages results in an overall average wage increase of 1.6% and widens the gender gap by 25.2%. By removing the 'AI premium' that women have in highly exposed AI jobs, the value of γ falls, negatively affecting the wages of female workers in occupations exposed to AI. Male workers benefit in this scenario, with an increase in their hourly wage of £0.6.

For our third counterfactual, we remove the female penalty in perceived productivity by setting κ_{women} to zero while keeping all other parameter estimates in θ unchanged. Under this setting, the average wage falls significantly by 10.1%, primarily due to a 17.7% decrease in male wages. The decline in female wages is minimal (0.5%), similar to the second counterfactual. This divergence in wage adjustments results in a 142.3% reduction in the gender gap, indicating that not only would the gap close, but women would, on average, earn higher wages than men.

For our fourth counterfactual, we modify the AIOE variable by capping its value at one. For occupations where the AIOE exceeds one, the value is substituted with one—meaning that occupations with AI exposure more than one standard deviation above the mean are now

treated as having the same level of exposure as those at the unit level. This imposition of an upper bound on AIOE leads to a 1.17% decrease in average wages, with the decline in male wages being 0.117 percentage points larger than that for female wages. In this case, the gender gap falls by 0.8%. Therefore, in the scenario where there is a limitation on AI such that the top-exposed occupations (those with an AIOE value one standard deviation above the mean) are affected, female workers would earn £0.18 less per hour and male workers would earn £0.23 less per hour.

7 Conclusion

In this paper, we have explored the potential effects of AI on the gender wage gap, focusing on its impact on job amenities and perceived productivity. Our findings indicate that women in occupations highly exposed to AI are perceived by firms to be less productive than their male counterparts by almost 4%. However, this productivity gap widens to 18.21% in jobs not exposed to AI, suggesting that the initial penalty in perceived productivity for women is largely, but not completely, mitigated in occupations with one standard deviation higher in AIOE. Additionally, our results indicate that while all workers face a reduction in job amenities value when working in intense AI jobs, this effect is larger for female workers. Moreover, the sign of our coefficents of interest (τ and ρ) seem to be alligned with the findings of Humlum and Vestergaard (2024) and Carvajal, Franco, and Isaksson (2024) in the surveys they run in Norway and Denmark respectively. Our simulated counterfactuals indicate that by eliminating the 'AI premium' that women experience in highly exposed AI jobs, the gender gap increases by more than 25%, driven by a significant increase in male wages. If occupations that are highly exposed to AI (those with an AIOE value greater than one) in the UK were to face a barrier causing their AIOE value to drop to one, there would be a slight decrease in the gender pay gap, but also a decline in the average hourly wages of all workers.

While the reasons behind these dynamics are still unclear, we have provided a novel perspective on the heterogeneous effects of AI on gender wage differentials, driven by its impact on job amenities and perceived productivity. In this study, we have focused our interest on the effects of AI on wages. However, our proposed methodology offers a significant opportunity to estimate counterfactual scenarios that examine the effects of AI not only on wages but also on matching probabilities. This is particularly valuable given the rapid evolution of AI and its uncertain impact on occupational distribution. This analysis represents an important policy tool for policymakers interested in studying changes in occupational sorting caused by

AI technology. For instance, STEM occupations are male-dominated, and the reasons behind this are not fully explained. It might be possible that male or female workers sort differently into occupations highly exposed to AI due to several factors, generating segregation in these types of jobs.

References

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

Aksoy, Cevat Giray, Berkay Özcan, and Julia Philipp (2021). "Robots and the gender pay gap in Europe". In: *European Economic Review* 134, p. 103693.

Albanesi, Stefania et al. (2023). *New technologies and jobs in Europe*. Tech. rep. National Bureau of Economic Research.

Autor, David and David Dorn (2013). "The growth of low-skill service jobs and the polarization of the US labor market". In: *American economic review* 103.5, pp. 1553–1597.

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

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

Becker, Gary S (2010). *The economics of discrimination*. University of Chicago press.

Bergmann, Barbara R (1974). "Occupational segregation, wages and profits when employers discriminate by race or sex". In: *Eastern economic journal* 1.2, pp. 103–110.

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

— (2020). "The gender pay gap: Have women gone as far as they can?" In: *Inequality in the United States*. Routledge, pp. 345–362.

Brussevich, Mariya, Ms Era Dabla-Norris, and Salma Khalid (2019). *Is technology widening the gender gap? Automation and the future of female employment*. International Monetary Fund.

Brynjolfsson, Erik, Tom Mitchell, and Daniel Rock (2018). "What can machines learn and what does it mean for occupations and the economy?" In: *AEA papers and proceedings*. Vol. 108. American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203, pp. 43–47.

Carvajal, Daniel, Catalina Franco, and Siri Isaksson (2024). "Will Artificial Intelligence Get in the Way of Achieving Gender Equality?" In: *NHH Dept. of Economics Discussion Paper* 03.

Dios-Murcia, Iria et al. (2021). "Trabajo en plataformas digitales:¿ una oportunidad para la oferta laboral femenina?" In: *ICE, Revista de Economıa* 921.

Dodge, Samuel and Lina Karam (2017). "A study and comparison of human and deep learning recognition performance under visual distortions". In: *2017 26th international conference on computer communication and networks (ICCCN)*. IEEE, pp. 1–7.

Dupuy, Arnaud and Alfred Galichon (2022). "A note on the estimation of job amenities and labor productivity". In: *Quantitative Economics* 13.1, pp. 153–177.

Felten, Manav Raj, and Seamans (2019). "The effect of artificial intelligence on human labor: An ability-based approach". In: *Academy of Management Proceedings*. Vol. 2019. 1. Academy of Management Briarcliff Manor, NY 10510, p. 15784.

— (2021). "Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses". In: *Strategic Management Journal* 42.12, pp. 2195–2217.

— (2023). "How will language modelers like chatgpt affect occupations and industries?" In: *arXiv preprint arXiv:2303.01157*.

Fernández-Macías, Enrique and Martina Bisello (2017). *Measuring the content and methods of work: A comprehensive task framework*. Tech. rep. Technical Report. European Foundation for the Improvement of Living and ...

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

Goldin, Claudia, Sari Pekkala Kerr, and Claudia Olivetti (2022). *When the kids grow up: women's employment and earnings across the family cycle*. Tech. rep. National Bureau of Economic Research.

Goos, Maarten, Alan Manning, and Anna Salomons (2008). "Recent changes in the European employment structure: The roles of technology, globalization and institutions". In: *Katholieke Universiteit Leuven*.

— (2014). "Explaining job polarization: Routine-biased technological change and offshoring". In: *American economic review* 104.8, pp. 2509–2526.

Gries, Thomas and Wim Naudé (2022). "Modelling artificial intelligence in economics". In: *Journal for labour market research* 56.1, p. 12.

Guarascio, Dario, Jelena Reljic, and Roman Stöllinger (2023). *Artificial Intelligence and Employment: A Look into the Crystal Ball*. Tech. rep. GLO Discussion Paper.

Hosny, Ahmed et al. (2018). "Artificial intelligence in radiology". In: *Nature Reviews Cancer* 18.8, pp. 500–510.

Humlum, Anders and Emilie Vestergaard (2024). "The Adoption of ChatGPT". In: *University of Chicago, Becker Friedman Institute for Economics Working Paper* 2024-50.

Martínez-Plumed, Fernando et al. (2020). "Does AI qualify for the job? A bidirectional model mapping labour and AI intensities". In: *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, pp. 94–100.

Mueller, Gerrit and Erik Plug (2006). "Estimating the effect of personality on male and female earnings". In: *Ilr Review* 60.1, pp. 3–22.

Stapleton, Katherine, Alexander Copestake, and Ashley Pople (2021). "AI, firms and wages: Evidence from India". In: *Available at SSRN 3957858*.

Tolan, Songül et al. (2021). "Measuring the occupational impact of ai: tasks, cognitive abilities and ai benchmarks". In: *Journal of Artificial Intelligence Research* 71, pp. 191–236.

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

Zarifhonarvar, Ali (2023). "Economics of chatgpt: A labor market view on the occupational impact of artificial intelligence". In: *Journal of Electronic Business & Digital Economics*.

A Log-likelihood gradients

The partial derivatives of the matching log-likelihood function are:

$$\frac{\partial \log \hat{L}_1(\theta)}{\partial A_k} = \frac{\partial \log \hat{L}_1(\theta)}{\partial \Gamma_k} = \sum_{i=1}^{n} \varphi_k(x_i, y_i) - n \sum_{i,j=1}^{n} \pi_{ij} \varphi_k(x_i, y_j) \tag{11}$$

$$\frac{\partial \log \hat{L}_1(\theta)}{\partial s^2} = \frac{\partial \log \hat{L}_1(\theta)}{\partial \sigma_1} = \frac{\partial \log \hat{L}_1(\theta)}{\partial \sigma_2} = \frac{\partial \log \hat{L}_1(\theta)}{\partial t} = 0 \tag{12}$$

The partial derivatives of the wage log-likelihood function with respect to all parameters but s^2 is:

$$\frac{\partial \log \hat{L}_2(\theta)}{\partial \theta_k} = s^{-2} \sum_{i=1}^{n} (W_i - w_i(\theta)) \frac{\partial w_i(\theta)}{\partial \theta_k} \tag{13}$$

and with respect s^2 is:

$$\frac{\partial \log \hat{L}_2(\theta)}{\partial s^2} = \sum_{i=1}^{n} \frac{(W_i - w_i(\theta))^2}{2s^4} - \frac{n}{2s^2} \tag{14}$$