

## **Artificial Intelligence Transparency Notice**

ADP uses artificial intelligence systems that include generative and traditional machine learning components. Specifically, in providing facts about your company or in surfacing answers about how to perform functions within the product, ADP relies on AI systems to help provide the narrative response, generate insights, make predictions, provide feedback, and provide initial drafts of certain documents (including job descriptions) for practitioner review. The AI system acts on ADP's extensive knowledge database, based on real interactions, as well as your company-specific data to develop these answers. The AI system is constrained to specific use cases in the types of questions it can answer. ADP uses rigorous methods to safeguard privacy and keep client data secure in a non-public environment. ADP provides ongoing human oversight to address data security, validity of outcomes, and protection against bias in all our AI systems. Validate any AI-generated content to ensure it is accurate, relevant, and complete for your company's purposes. ADP notes that not all companies use AI offered by ADP and, therefore, the information in this notice may not be relevant to the specific company depending on their use.

### **Candidate Relevancy FAQ's**

*Employers who post jobs on ADP's recruiting platforms may refer to an applicant's Candidate Relevancy or Profile Relevance score. Candidate Relevancy and Profile Relevance rely on artificial intelligence and machine learning to provide an initial comparison of an applicant's education, experience, and skills against the education, experience, and skills requirements in the job description. This is intended to be one of many factors that a potential employer will review in making its interview decisions; there are no cut-off scores and all applications remain visible to employers. Candidates who opt out will have their score listed as "Not available."*

*These FAQs provide additional information about the data these tools collect, store, and retain, and the results of the most recent impartial evaluations of these tools.*

#### **1. What is Candidate Relevancy?**

ADP's Candidate Relevancy and Profile Relevance tools (for ease of reference both will jointly be referred to as "Candidate Relevancy" unless otherwise noted) use artificial intelligence and machine learning algorithms to conduct an initial review of an application, and are designed to be utilized by employers as one tool, among others, in the hiring process.<sup>1</sup> Specifically, Candidate Relevancy conducts a mathematical assessment of how close the skills, education and/or experience on an applicant's resume match the skills, education, and/or experience listed on the relevant job description. This process quantifies the "relevance" between the applicant's resume

---

<sup>1</sup> The Candidate Relevancy score is displayed to employers using ADP's Recruitment Management product, while the Profile Relevance score is displayed to employers using ADP's WorkforceNow Recruitment platform.

and the job posting. The Candidate Relevancy model also leverages past decisions derived from millions of resumes and job descriptions where the selection decision is already known.

The scores are intended to be used as one of many factors by an employer in determining who to advance to the next round in the hiring process. Candidate Relevancy is not intended to replace human judgment during any step of the recruitment process and is designed in such a way that there are no cut-off scores that would eliminate applicants from being visible to employers in the user interface. Employers are provided access to all applications, enabling them to make human decisions on which candidates to pursue.

## **2. How is the Candidate Relevancy score determined?**

The Candidate Relevancy model first parses the information concerning the education, experience, and skills contained in the applicant's resume or application and in the relevant job description. This information is formatted to allow a mathematical assessment to be conducted of how close the applicant's education, skills, and experience match those found in the relevant job description. Candidate Relevancy does not extract or utilize the applicant's name, address, race, ethnicity, gender or protected demographic information.

Each job requisition is classified using a job and sector taxonomy. The Candidate Relevancy model creates three sub-scores indicating how close the applicant's education, skills, and experience matches those found in the job description. The three scores are then weighted to create the Candidate Relevancy Score. The weights sum to 1 and reflect the relative importance of each component. Since the job descriptions do not define the importance of each component, the importance (i.e., the weights) must be estimated empirically from the data. Separate weights are created for each sector in which the open job resides. The weights are determined by a machine learning model.

The resulting weighted score (the final Candidate Relevancy score) is intended to be used by an employer as only one tool, among others, to aid in the selection of whom to interview or prioritize during the hiring pipeline.

## **3. What data does Candidate Relevancy collect and what are ADP's retention policies regarding the information?**

Type of Data	Collected from	Retention Policy
Resume data	ADP Workforce Now Recruitment or ADP Recruitment Management	Three years
Job descriptions	ADP Workforce Now Recruitment or ADP Recruitment Management	Three years

#### **4. Is Candidate Relevancy an automated employment decision tool covered by New York City Local Law 144 ("the NYC Ordinance")?**

The NYC Ordinance covers automated screening or selection tools that provide "output"—such as scores, classifications, or recommendations—to an employer, and which are used to significantly assist or substitute a human's decision-making process. Under the NYC Ordinance, to substantially assist or substitute a human's decision-making process means: (1) to rely solely on a simplified output without consideration to other factors; (2) to use a simplified output as a consideration in a list of criteria but weight the output more heavily than other criteria the set; or (3) to use the output to overrule human decision-making conclusions.

Candidate Relevancy is not intended by ADP to be relied upon solely by employers in making employment decisions and is not meant to substantially assist or replace discretionary decision making in employment decisions. Moreover, Candidate Relevancy is not intended to be used as a criterion that is weighted more than any other criterion in making employment decisions and is not intended to be used to overrule conclusions derived from other factors, including human decision-making.

Candidate Relevancy is intended to be one source of assistance in helping to prioritize candidates selected for next steps. Education, skills, and experience must be evaluated and validated by employers through person-to-person interviews and background checks, among other things. Candidate Relevancy is not intended to replace human judgment during any step of the recruitment process and is designed in such a way that there are no cut-off scores that would eliminate candidates from being visible to employers in the user interface. Employers are thereby provided access to all candidates, enabling them to make human decisions on which candidates to pursue.

If Candidate Relevancy is used as intended by ADP, ADP does not believe Candidate Relevancy to be an automated employment decision tool as defined by the New York City Ordinance and its related final rules.

Nothing herein is intended to be a legal opinion and does not constitute legal advice. You should consult with an attorney before taking any action in reliance on the information provided herein including whether Candidate Relevancy is an automated employment decision tool.

#### **5. Did ADP conduct a bias audit on Candidate Relevancy?**

Yes. At ADP integrity is everything and is at the foundation of how we design and develop our solutions and services. Although ADP believes that Candidate Relevancy, if used as intended by ADP, does not fall within the scope of the NYC Ordinance, ADP is committed to ensuring that transparency and accountability is embedded in ADP's offerings.

ADP obtained an independent bias audit of Candidate Relevancy and Profile Relevance from BLDS, LLC, an independent auditor, in April of 2024. The independent auditors concluded that no

valid statistical evidence of bias is present in the scoring produced by Candidate Relevancy or Profile Relevance.

## 6. What was the result of the bias audit conducted on Candidate Relevancy?

In April of 2024, an independent auditor, BLDS, LLC, performed an impartial evaluation of Candidate Relevancy. The independent auditors concluded that no valid statistical evidence of bias is present.

A summary of the scoring rates and impact ratios<sup>2</sup> based on sex and race/ethnicity and the intersection of sex and race/ethnicity, and adjusted for Simpson's Paradox, are set forth in the following charts:

<b>Sex Categories</b>			
	<b>Applicants</b>	<b>Scoring Rate</b>	<b>Impact Ratio</b>
<b>Female</b>	1,030,417	49.6%	1.000
<b>Male</b>	868,162	48.5%	0.979
Unknown Gender	1,838,419	--	--

<b>Race/Ethnicity Categories</b>			
	<b>Applicants</b>	<b>Scoring Rate</b>	<b>Impact Ratio</b>
<b>Asian</b>	233,768	45.6%	0.874
<b>Black or African American</b>	452,625	48.9%	0.938
<b>Hispanic or Latino</b>	320,000	49.7%	0.954
<b>Two or More Races</b>	72,612	50.4%	0.966
<b>White</b>	716,986	52.2%	1.000
Unknown Race/Ethnicity	1,948,813	--	--

<b>Intersectional Categories</b>				
		<b>Applicants</b>	<b>Scoring Rate</b>	<b>Impact Ratio</b>
<b>Female</b>	<b>Asian</b>	98,422	47.0%	0.905
	<b>Black or African American</b>	278,254	48.7%	0.937
	<b>Hispanic</b>	159,439	49.4%	0.950
	<b>Two or More Races</b>	39,307	50.3%	0.968
	<b>White</b>	368,641	52.0%	1.000
<b>Male</b>	<b>Asian</b>	125,704	43.7%	0.840
	<b>Black or African American</b>	159,754	47.7%	0.919

<sup>2</sup> Consistent with the New York City Ordinance, impact ratio means either (1) the selection rate for a category divided by the selection rate of the most selected category or (2) the scoring rate for a category divided by the scoring rate for the highest scoring category.

	<b>Hispanic</b>	145,714	48.9%	0.941
	<b>Two or More Races</b>	25,232	49.2%	0.947
	<b>White</b>	346,179	50.5%	0.972
Unknown Intersectionality		1,988,994	--	--

American Indian or Alaska Natives or the Native Hawaiian or Other Pacific Islanders were not included in computing the Impact Ratio because both categories had less than 1% of the population and the New York City Ordinance does not require their inclusion when computing the Impact Ratio. In the opinion of the independent auditors, the inclusion of such small numbers would allow the race/ethnicity or intersectional categories of American Indian or Alaska Natives or Native Hawaiian or Other Pacific Islanders to be the highest selection rate based on a small number of cases. Allowing such a small sample as the reference group to judge other categories is questionable as the standard for judging the results of other categories for many jobs/sectors would be set based on only a handful of cases. The table below reports the data adjusted for Simpson’s Paradox on the categories that were not used in computing the Impact Ratio.

<b>Populations Less Than 1%</b>		
	<b>Applicants</b>	<b>Scoring Rate</b>
<b>Native American / Alaska Native</b>	6,382	48.1%
<b>Native Hawaiian / Pacific Islander</b>	4,667	50.8%
<b>Female Native American / Alaska Native</b>	3,508	48.5%
<b>Male Native American / Alaska Native</b>	2,276	48.6%
<b>Female Native Hawaiian / Pacific Islander</b>	2,073	45.8%
<b>Male Native Hawaiian / Pacific Islander</b>	1,765	50.8%

This analysis was conducted across all uses of Candidate Relevancy where sufficient self-ID information was available. Nothing in these FAQ’s should be taken as a guarantee that a particular client’s use of Candidate Relevancy will never result in adverse impact or bias.

**7. What was the result of the bias audit conducted on Profile Relevance?**

An independent bias audit of Profile Relevance was also conducted by BLDS, LLC in April of 2024.<sup>3</sup> The independent auditors concluded that no valid statistical evidence of bias is present.

This analysis defined “selection” as candidates placed in the “High” category and in the “High or Medium” category. A summary of the selection rates and impact ratios based on sex and race/ethnicity and the intersection of sex and race/ethnicity, and adjusted for Simpson’s Paradox, are set forth in the following charts:

---

<sup>3</sup> Candidate Relevancy and Profile Relevance rely on the same algorithm to produce a numerical relevancy score (1 to 100). Candidate Relevancy displays the numerical score (1 to 100) to recruiters, while Profile Relevance converts the numerical score into a High, Medium, or Low relevancy category. Because the interface is different at this time, ADP obtained separate independent bias audits for Candidate Relevancy and Profile Relevance.

<b>Sex Categories</b>				
<b>Selection Classified as High</b>				
	<b>Applicants</b>	<b>Selections</b>	<b>Scoring Rate</b>	<b>Impact Ratio</b>
<b>Female</b>	5,633,755	2,285,051	40.6%	1.000
<b>Male</b>	4,667,322	1,874,397	40.2%	0.990

<b>Selection Classified as High or Medium</b>				
	<b>Applicants</b>	<b>Selections</b>	<b>Scoring Rate</b>	<b>Impact Ratio</b>
<b>Female</b>	5,767,615	4,267,458	74.0%	1.000
<b>Male</b>	4,798,518	3,536,508	73.7%	0.996
Unknown Sex	4,031,410	--	--	--

<b>Race / Ethnicity Categories</b>				
<b>Selection Classified as High</b>				
	<b>Applicants</b>	<b>Selections</b>	<b>Scoring Rate</b>	<b>Impact Ratio</b>
<b>Asian</b>	692,402	268,583	38.8%	0.934
<b>Black or African American</b>	2,374,766	969,379	40.8%	0.983
<b>Hispanic or Latino</b>	1,646,306	678,113	41.2%	0.992
<b>Two or More Races</b>	371,327	154,249	41.5%	1.000
<b>White</b>	3,587,705	1,470,600	41.0%	0.987

<b>Selection Classified as High or Medium</b>				
	<b>Applicants</b>	<b>Selections</b>	<b>Scoring Rate</b>	<b>Impact Ratio</b>
<b>Asian</b>	718,638	524,534	73.0%	0.972
<b>Black or African American</b>	2,414,565	1,795,712	74.4%	0.990
<b>Hispanic or Latino</b>	1,676,917	1,253,999	74.8%	0.996
<b>Two or More Races</b>	375,123	281,717	75.1%	1.000
<b>White</b>	3,683,029	2,754,906	74.8%	0.996
Unknown Race/Ethnicity	6,116,411	--	--	--

<b>Intersectional Categories</b>					
<b>Selection Classified as High</b>					
		<b>Applicants</b>	<b>Selections</b>	<b>Scoring Rate</b>	<b>Impact Ratio</b>
<b>Female</b>	<b>Asian</b>	302,583	118,824	39.3%	0.926
	<b>Black or African American</b>	946,081	384,014	40.6%	0.957
	<b>Hispanic/Latino</b>	861,334	356,765	41.4%	0.977
	<b>Two or More Races</b>	207,645	88,041	42.4%	1.000
	<b>White</b>	1,885,642	770,850	40.9%	0.964
<b>Male</b>	<b>Asian</b>	357,504	136,710	38.2%	0.902

	<b>Black or African American</b>	946,081	384,014	40.6%	0.957
	<b>Hispanic/Latino</b>	730,545	299,231	41.0%	0.966
	<b>Two or More Races</b>	133,167	56,023	42.1%	0.992
	<b>White</b>	1,637,864	661,369	40.4%	0.952

<b>Selection Classified as High or Medium</b>					
		<b>Applicants</b>	<b>Selections</b>	<b>Scoring Rate</b>	<b>Impact Ratio</b>
<b>Female</b>	<b>Asian</b>	312,571	229,396	73.4%	0.933
	<b>Black or African American</b>	1,395,522	1,040,362	74.6%	0.963
	<b>Hispanic/Latino</b>	874,174	653,795	74.8%	0.954
	<b>Two or More Races</b>	209,140	157,712	75.4%	0.984
	<b>White</b>	1,932,941	1,440,814	74.5%	0.950
<b>Male</b>	<b>Asian</b>	370,579	268,744	72.5%	0.928
	<b>Black or African American</b>	962,302	714,221	74.2%	0.961
	<b>Hispanic/Latino</b>	743,116	554,810	74.7%	0.958
	<b>Two or More Races</b>	134,099	101,607	75.8%	1.00
	<b>White</b>	1,688,353	1,254,615	74.3%	0.947
Unknown Intersectional		6,286,786	--	--	--

American Indian or Alaska Natives or the Native Hawaiian or Other Pacific Islanders were not included in computing the Impact Ratio because both categories had less than 1% of the population, and the New York City Ordinance does not require their inclusion when computing the Impact Ratio. In the opinion of the independent auditors, the inclusion of such small numbers would allow the race/ethnicity or intersectional categories of American Indian or Alaska Natives or Native Hawaiian or Other Pacific Islanders to be the highest selection rate based on a trivial number of cases. Allowing such a small sample as the reference group to judge other categories is questionable as the standard for judging the results of other categories for many jobs/sectors would be set based on only a handful of cases. The table below reports the data, adjusted for Simpson's Paradox, on the categories that were not used in computing the Impact Ratio.

<b>Populations Less Than 1%</b>			
<b>Selection Classified as High</b>			
	<b>Applicants</b>	<b>Selections</b>	<b>Selection Rate</b>
<b>Native American / Alaska Native</b>	44,790	20,129	44.9%
<b>Native Hawaiian / Pacific Islander</b>	26,195	12,309	47.0%
<b>Female Native American / Alaska Native</b>	22,379	10,263	45.9%
<b>Male Native American / Alaska Native</b>	15,442	7,505	48.6%
<b>Female Native Hawaiian / Pacific Islander</b>	12,875	6,314	49.0%
<b>Male Native Hawaiian / Pacific Islander</b>	8,963	4,647	51.9%

<b>Selection Classified as High or Medium</b>			
	<b>Applicants</b>	<b>Selections</b>	<b>Selection Rate</b>
<b>Native American / Alaska Native</b>	44,865	34,595	77.1%
<b>Native Hawaiian / Pacific Islander</b>	26,214	20,793	79.3%
<b>Female Native American / Alaska Native</b>	22,415	17,349	77.4%
<b>Male Native American / Alaska Native</b>	15,442	12,210	79.1%
<b>Female Native Hawaiian / Pacific Islander</b>	12,903	10,319	80.0%
<b>Male Native Hawaiian / Pacific Islander</b>	8,978	7,296	81.3%

This analysis was conducted across all uses of Profile Relevance where sufficient self-ID information was available. Nothing in these FAQ's should be taken as a guarantee that a particular client's use of Profile Relevance will never result in adverse impact or bias.

**8. Can applicants opt out of having their resume reviewed by Candidate Relevancy? What happens if someone opts out?**

All applicants are included in the applicant queue for a recruiter to review. Individuals applying through ADP's recruiting platforms can choose not to have their application reviewed by Candidate Relevancy or Profile Relevance tools. Each opt-out choice is job-specific and opts the candidate out for the specific job posting only. For applicants who have chosen to opt out, their score will be listed as "Not Available," which is the same indicator used if a relevancy score is



unavailable for reasons other than opt-out (e.g., technical issues, poor resolution on resume pdf, etc.).

### **ADP's Commitment to Ethical Artificial Intelligence**

For more information about ADP's commitment to ethical artificial intelligence please refer to <https://www.adp.com/about-adp/artificial-intelligence.aspx>.

For any questions or inquiries, please contact [AIEthics@adp.com](mailto:AIEthics@adp.com).

This document and all of its contents is the property of ADP, Inc. This document is for information purposes only. Diagrams, tables, percentages and/or outcomes used in this document is for illustration purposes only. Individual outcomes vary by customer. ADP's customers are solely responsible for its use of ADP technology. ADP will not be responsible for any liability, loss or damage of any kind resulting from or connected with the use of this document.