

# ► AI in human resource management

## The limits of empiricism

Authors / Janine Berg, Hannah Johnston





Attribution 4.0 International (CC BY 4.0)

This work is licensed under the Creative Commons Attribution 4.0 International. See: <https://creativecommons.org/licenses/by/4.0/>. The user is allowed to reuse, share (copy and redistribute), adapt (remix, transform and build upon the original work) as detailed in the licence. The user must clearly credit the ILO as the source of the material and indicate if changes were made to the original content. Use of the emblem, name and logo of the ILO is not permitted in connection with translations, adaptations or other derivative works.

**Attribution** – The user must indicate if changes were made and must cite the work as follows: Berg, J., Johnston, H. *AI in human resource management: The limits of empiricism*. ILO Working Paper 154. Geneva: International Labour Office, 2025. © ILO.

**Translations** – In case of a translation of this work, the following disclaimer must be added along with the attribution: *This is a translation of a copyrighted work of the International Labour Organization (ILO). This translation has not been prepared, reviewed or endorsed by the ILO and should not be considered an official ILO translation. The ILO disclaims all responsibility for its content and accuracy. Responsibility rests solely with the author(s) of the translation.*

**Adaptations** – In case of an adaptation of this work, the following disclaimer must be added along with the attribution: *This is an adaptation of a copyrighted work of the International Labour Organization (ILO). This adaptation has not been prepared, reviewed or endorsed by the ILO and should not be considered an official ILO adaptation. The ILO disclaims all responsibility for its content and accuracy. Responsibility rests solely with the author(s) of the adaptation.*

**Third-party materials** – This Creative Commons licence does not apply to non-ILO copyright materials included in this publication. If the material is attributed to a third party, the user of such material is solely responsible for clearing the rights with the rights holder and for any claims of infringement.

Any dispute arising under this licence that cannot be settled amicably shall be referred to arbitration in accordance with the Arbitration Rules of the United Nations Commission on International Trade Law (UNCITRAL). The parties shall be bound by any arbitration award rendered as a result of such arbitration as the final adjudication of such a dispute.

For details on rights and licensing, contact: [rights@ilo.org](mailto:rights@ilo.org). For details on ILO publications and digital products, visit: [www.ilo.org/publns](http://www.ilo.org/publns).

---

ISBN 9789220428610 (print), ISBN 9789220428627 (web PDF), ISBN 9789220428634 (epub), ISBN 9789220428641 (html). ISSN 2708-3438 (print), ISSN 2708-3446 (digital)

<https://doi.org/10.54394/NMSH7611>

---

The designations employed in ILO publications, which are in conformity with United Nations practice, and the presentation of material therein do not imply the expression of any opinion whatsoever on the part of the ILO concerning the legal status of any country, area or territory

or of its authorities, or concerning the delimitation of its frontiers or boundaries. See: [www.ilo.org/disclaimer](http://www.ilo.org/disclaimer).

The opinions and views expressed in this publication are those of the author(s) and do not necessarily reflect the opinions, views or policies of the ILO.

Reference to names of firms and commercial products and processes does not imply their endorsement by the ILO, and any failure to mention a particular firm, commercial product or process is not a sign of disapproval.

Information on ILO publications and digital products can be found at: [www.ilo.org/research-and-publications](http://www.ilo.org/research-and-publications)

ILO Working Papers summarize the results of ILO research in progress, and seek to stimulate discussion of a range of issues related to the world of work. Comments on this ILO Working Paper are welcome and can be sent to [berg@ilo.org](mailto:berg@ilo.org).

Authorization for publication: Caroline Fredrickson, Director, Research Department

ILO Working Papers can be found at: [www.ilo.org/research-and-publications/working-papers](http://www.ilo.org/research-and-publications/working-papers)

**Suggested citation:**

Berg, J., Johnston, H. 2025. *AI in human resource management: The limits of empiricism*, ILO Working Paper 154 (Geneva, ILO). <https://doi.org/10.54394/NMSH7611>

## Abstract

---

The rapid integration of artificial intelligence (AI) into Human Resource Management (HRM) is transforming how organizations recruit, manage, and evaluate their workforces. While proponents champion AI as a means to enhance efficiency, reduce bias, and align HR practices with strategic business goals, this paper argues that such optimism is misplaced. Drawing on a critical review of AI's application across four core HRM functions—recruitment, compensation, scheduling, and performance management—this paper identifies significant risks and limitations arising from the fundamental structure of AI systems.

Central to the analysis is a three-parameter framework for assessing AI tools: their *objective*, the *data* they rely upon, and how they are *programmed*. The paper shows that across HR functions, AI systems frequently operationalize reductive or poorly aligned objectives, rely on low-quality or biased data, and are programmed in non-transparent ways that undermine their usefulness. These structural shortcomings not only undermine the effectiveness of AI systems but also introduce legal, ethical, and practical risks for firms and their workers.

**Keywords:** artificial intelligence, human resource management, data analytics, algorithmic management

## About the authors

---

**Janine Berg** is Senior Economist and Head of the Effective Labour Institutions Unit in the Research Department of the ILO. Since joining the ILO in 2002, she has conducted research on the economic and social effects of labour laws as well as provided technical assistance on policies for generating jobs and improving working conditions. She is the author of several books and numerous articles on employment, labour market institutions and the digital transformation of work.

**Hannah Johnston** is an Assistant Professor in the School of Human Resources Management at York University in Toronto, Canada, specializing on the digitalization of work. Prior to joining York, Hannah was a postdoctoral fellow at Northeastern University in Boston and also worked professionally at the International Labour Organization. Hannah has a longstanding interest in the platform economy and is a collaborator with Oxford University's Fairwork Project. Her recent publications can be found in journals including *Industrial and Labor Relations Review*, *Work and Occupations*, and the *International Labour Review*.

## Table of contents

---

Abstract	01
About the authors	01
<hr/>	
► Introduction	04
<hr/>	
► 1 How did we get here? The rise of “people analytics” and managing workers through data	05
<hr/>	
► 2 The workings of AI systems: Objective, data and programming	08
Recruitment	09
Sourcing	10
Screening, Interviewing and Selection	11
Compensation	13
Scheduling work	16
Performance management	19
<hr/>	
► 3 AI in HRM: Unbridled optimism	23
<hr/>	
► 4 What is an HR manager to do with AI?	26
<hr/>	
► Conclusion	28
<hr/>	
References	29
Acknowledgements	36

## List of Figures

---

<b>Figure 1. McKinsey infographic on the benefits of people analytics</b>	<b>06</b>
<b>Figure 2. The workings of AI systems</b>	<b>09</b>

## ► Introduction

---

Motivated by a desire to more efficiently and effectively manage people in organizations, HR managers are using the programming and analytical capacities of AI to fulfill key HR functions – including selection and recruitment of personnel, compensation determination and structure, performance review and evaluation, and the organization of working time. Largely absent from the rush to integrate AI, however, has been a comprehensive assessment of whether, and under what circumstances, AI is useful for the management of people within organizations. Yet despite the lack of assessment, the ‘AI for HR’ industry and the adoption of these tools and systems by individual firms and organizations is burgeoning.

This paper presents a framework for understanding and evaluating the potential benefits and possible risks or harms presented by AI systems in workforce management. Following a section of the paper documenting the historical context of the Human Resources field that has given rise, first to people analytics and then to AI, the paper presents a framework based on three inter-related parameters that can help assess the quality, legality, and suitability of AI systems used in the field. These are: (1) the system objective, (2) the data it is built on and relies on, and (3) how the AI system is programmed. Drawing on existing literature about how AI is being used for workforce management, the paper applies the three-parameter framework to map the contours of AI use relative to four key Human Resource Management functions where adoption of AI technologies has been prominent: recruitment, compensation, scheduling, and performance management. Our discussion section, “Unbridled optimism” views the disciplinary and occupational history of HR alongside the findings on HR’s emergent use of AI. We argue that the search for occupational legitimacy by HR professionals has fostered a preoccupation with numeracy and positivism that has provided fertile ground for the ‘evidence-based solutions’ that AI systems purport to offer. This tendency towards positivism, we contend, is likely to result in widespread adoption of AI under circumstances that create risks for workers, liabilities for firms, and costs for society. Given the seemingly inevitable transformation of work due to AI, we argue for the need for HR professionals to improve their understanding of the workings of AI systems so that they can better judge their potential and limitations, and that this is best achieved when they participate in the design of systems that are implemented in their workplaces.

## ► 1 How did we get here? The rise of “people analytics” and managing workers through data

---

Human Resource Management emerged in the 1950s as a distinct field of study and practice concerned with the management of people in organizations. Carved out from the broader discipline of Industrial Relations, which has historically examined labour relations in the context of unions and collective bargaining – or collective employment relations – Human Resource Management is most associated with an employer’s relations with individual employees (Kaufman 2001). As practitioners, HR managers are typically engaged in facilitating a range of personnel functions – including recruitment and selection, compensation, scheduling and performance management and promotion – to achieve the organization’s goals.

Since the 1980s, the dominant paradigm guiding HR managers has been ‘strategic HRM’ (Paauwe and Boon 2018). Building on HR’s origins and the foundational concept of scientific management, strategic HRM aims to link firm performance to the specific methods and practices deployed within the firm to manage their workers. When first introduced, this approach represented a significant shift in firm management. Prior to this shift, organizational ‘strategy’ referred to a firm’s perspective or ‘world view’, its intended plan and patterned behaviour, as well as a firm’s use of ploys to outwit competitors (Mintzberg 1987). However, in the 1980s firm strategy began to focus more on causality, with researchers and practitioners attempting to identify inputs (e.g. management practices) and outputs (e.g. market performance), quantify them, and derive causal connections between the two. Operationalizing this strategic approach, as many have argued, has largely been “geared towards specific numerical targets” (Wood and Kispál-Vitai 2017). Firms seek not only to derive general conclusions or trends through quantification, but also to compare a range of strategies for the purpose of distinguishing a singular ‘best practice’.

This epistemological shift has fueled the development of new metrics and datapoints that can be used to analyze the relationship between workforce management and firm performance. In turn, HR managers can use these new sources of data about the workforce to inform decision-making. This type of Evidence-Based Management technique (EBM) has been largely promoted in strategic HRM to help overcome pitfalls, such as relying on personal experience or managerial whims, or the tendency to mimic the strategies or approaches of top performers (Pfeffer and Sutton 2006; Rousseau 2006; Reay, Berta, and Kohn 2009). The result has been a vast arsenal of digitally enabled workplace and work-related tools that capture, collect and analyze worker behavior and performance data. Since the 1980s, this has driven the growing field of people analytics – defined as “the use of measurement and analysis techniques to understand and optimize the people side of business” (Enderes and Shannon 2019) – and ultimately has paved the path towards the adoption of AI for HRM.

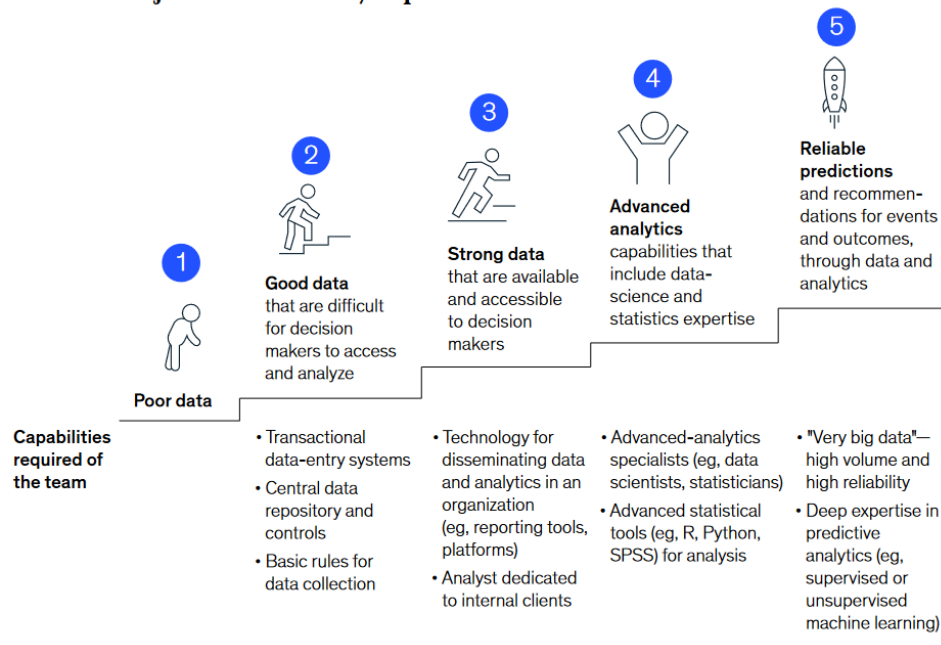
Within the field of people analytics and the development of AI, more data and better data are commonly viewed as precursors for robust and powerful systems. People analytics data may pertain to workers’ demographics, descriptive information about the location or nature of the job, performance, training or professional history, or tenure. While every organization has data on its workforce, advanced analysis for workforce optimization and planning can only be achieved when data are high quality, robust and plentiful, and when organizations have adequately trained staff to process and make sense of them. Figure 1 is an infographic from a McKinsey publication on the “virtue” of people analytics, stressing that the power of such systems to provide workforce insights increases with the volume and quality of workforce data (Ledet et al. 2020).



► Figure 1. McKinsey infographic on the benefits of people analytics

Exhibit 1

**The best people analytics teams may take one step back for every two steps up, but their trajectories are always upward.**



Source: (Ledet et al. 2020).

When first introduced, HR professionals used people analytics to assess patterns about their workforce, informing a wide range of HR functions including recruitment and hiring, promotion, compensation, and health and safety (Giermindl et al. 2022). With more and higher-quality “big data” and increases in computing power, people analytics is being propelled from correlation analysis into the world of pattern-based prediction, and Human Resource professionals, once responsible for executing a wide range of functions, are, in some instances, relinquishing these responsibilities to algorithms and AI. This decades-long strategic shift provides critical context for understanding why the field of HRM has embraced the use of AI; the particular harms that may emerge from the use of AI; and why the field of HR is largely blind to them.

AI is distinguished by vast quantities of data and rapid quantitative analysis. These new sources of data have held particular appeal to a profession that has oriented itself towards EBM. Additionally, AI is popularly portrayed as innovative and cutting edge, and the use of AI technologies and tools is thus seen as a way to elevate the HR profession. This motivation is also underpinned by important context as historically many have tended to regard HR professionals as administrative functionaries, who – lacking power – are engaged in mere bureaucratic service delivery without adding any real value to the organization (Wright 2008; Legge 1978).

Although the strategic focus of HRM in the 1980s (along with other activities such as the formation of professional associations and educational and training courses) provided one avenue for the occupation to recast itself on equal footing with other managerial professions (Legge 1978; Cayrat and Boxall 2023), critics maintain that the field has still not provided evidence of its worth (Wright 2008; Alvesson 2008; Kryscynski et al. 2018; Cayrat and Boxall 2023; Hammonds 2005). Thus, this persistent quest for occupational legitimacy presents another important motivation for HR’s embrace of people analytics and AI: by embracing evidence-based management (EBM),

and the data accumulating and data-intensive tools and technologies that facilitate it, HR practitioners might provide evidence of their value to the firm.

This approach, however, has not been without criticism. Scholars argue that while empirical research linking HR practices to organizational performance can at times demonstrate a clear association, the relationship between these variables is under-theorized (Fleetwood and Hesketh 2008). The empirical approach of HR scholars has proven adequate for generating predictive dimensions of theory rooted in prior observations of ‘what’ is happening and ‘how’; however, robust theories should also be capable of explaining the mechanisms and reasoning behind the mechanisms and causal relationships (Guest 2025; Fleetwood and Hesketh 2008). It is with respect to this latter explanatory dimension that HRM theory has fallen short. Without an adequate or clearly articulated theory, research will “also lack an adequate rationale for the choice of phenomena that will eventually become the variables” used to drive future predictive theorization (ibid); 127). Under these conditions, rather than explaining why organizations succeed or fail and what other systemic conditions or discrete practices may contribute to this outcome (Guest 2025), HR managers and researchers alike are prone to replicate past practice simply because such practices have been previously examined. HR theories, in turn, remain poorly articulated and lack conceptual clarity (Boon et al. 2019). This risk of ‘measurement without theory’ is amplified in the context of AI, which by definition is not about building an explicit theory, but rather about building patterns through big data (Elragal and Klischewski 2017).

Proponents of AI integration, including many global consulting firms, promise exactly this: added value from the use of their newly developed technologies and tools – and specifically, that the predictive capacity of AI will help to reveal the recipe for an effective allocation of human resources. Companies are rapidly obliging. ISG, an organizational change management company, found in 2023 that one out of every three organizations was prioritizing AI and analytics in their HR and technology strategy. This finding was based on a survey of enterprise leaders at firms employing between 5,000 and 50,000 workers. As the report explained, “leading HR technology providers, such as such as Oracle, [IBM], SAP SuccessFactors and Workday, are focused on embedding AI, machine learning and analytics in the core platform” (ISG 2023). Indeed, most of the software is provided through third-parties and is typically an add-on to other services already provided. This facilitates the work of HR professionals who often lack adequate training in what kinds of analytical questions to ask or how to interpret quantitative findings (Giermendl et al. 2022; Kryscynski et al. 2018).

Yet despite high rates of technological integration, including AI, the same ISG survey found that less than half of organizations surveyed realized business value from their investments. What has emerged is thus a paradox: although troves of data and developments in the field of AI promise to provide HR professionals with the necessary data and analysis to help them effectively and efficiently allocate human resources, these technologies have not yet delivered much value to firms and organizations. Why not? What does the actual evidence on the implementation of AI in HRM bear? Does the design of these AI systems allow them to deliver as promised?

## ► 2 The workings of AI systems: Objective, data and programming

---


To assess how effective AI has been for realizing the overarching goals of HR, it is necessary to delineate what these goals are, and to unpack the workings of AI system to better understand potential pitfalls. AI systems are composed of three inter-related parameters: (1) the system objective, (2) the data it is built on and relies on, and (3) how it is programmed (See figure 2). The quality of each of these parameters differentiates systems that work well from those that do not.

Beginning with the system objective, while defining such an aim may appear straightforward, how these aims are operationalized is more complicated. When AI systems are developed with neutral aims – such as determining the shortest route between locations – it is easy to rely on identifiable and relevant variables and the findings are easy to interpret. But most human resource functions involve the “fleshy, messy, indeterminate stuff of everyday life” – stuff that are difficult to capture in a discrete variable (Katz 2001). This can, as described by Sandy Gould (Gould 2024), create practical challenges, which she explains as the need to, “accept resource constraints on what is measured, as well as the ontological limitations (some things are never going to be open to ‘direct’ measurement by any conceivable means)” (106).

Data, meanwhile, are a necessary input to AI systems. Limitations in data have been the focus of much critique, with two particularly salient concerns. First, there is the question of data quality. AI systems rely on training data to ‘learn’ the connections and patterns that provide the foundation upon which decisions are made (Whang et al. 2023). When these data are poor quality, AI systems yield poor quality outputs. As the saying goes: ‘garbage in, garbage out’. A second issue concerns the suitability of the data. In bespoke AI systems that are developed internally by or for a singular firm, training data may be internal to the firm and consist of past operations (Kresge 2020). When systems are developed for ‘off the shelf’ use, they rely on more generalized datasets that can either be purchased via a growing data market (Zuboff 2019) or otherwise compiled from available data that are deemed to be relevant sources for the system (Muldoon et al. 2023). The appropriateness of more generalized datasets ought to be evaluated on a case-by-case basis but rarely is. Equally, when issues stemming from non-representative data arise, a commonly suggested solution is that *more data* or *more representative* data are needed. However, these types of problems also raise questions about the suitability of the data to meet the stated objective.

Those cautious of AI have also pointed to the potential problems associated with the third parameter: programming. Algorithms are a key decision-making feature at the core of AI systems. They can be defined, in their most basic sense, as a set of rules executed through computer programming code with a particular aim or objective. Algorithms function with varied levels of autonomy and human involvement; these characteristics are also determined by their design and deployment (OECD 2021). The most basic algorithms merely execute a list of prescribed instructions or offer data insights or recommendations that can assist human decision making; basic algorithms only change with intervention from those who use or program them. In these instances, bias can be intentionally or unintentionally embedded in the computer code by the programmer. Issues also arise with more autonomous ‘intelligent’ ‘machine learning’ algorithms like those associated with and used in AI systems. Programming can become fraught as the system ‘evolves’. Self-learning systems, although still developed and deployed by humans, often operate in opaque ways and their precise functionalities may elude even those who built them.

► Figure 2. The workings of AI systems



Source: Authors' elaboration.

For those who warn of the detrimental impact that AI could have in the world of work, how these three parameters of AI systems (objective, data, and programming) are chosen and deployed could jeopardize work quality and workplace fairness.

Moreover, the use of third parties' 'off-the-shelf' systems can heighten risks. First, the algorithms embedded in AI systems may be poorly understood by firms and workers using them because of intellectual property concerns; and firms using externally developed systems may have limited ability to modify how the system functions to meet their specific needs. Second, deploying these types of technologies in contexts where real-life data differ from training data can lead to discrimination or poor outcomes. Third, concerns about workers' data privacy have also emerged. While firms have long tracked a host of data about worker performance, the use of third-party systems often involves complex licensing arrangements that may include provisions that require workers' data to be shared with the AI developer. Such sharing raises important questions about consent, privacy, and even the commodification of workers' data (Kresge 2020).

Within purpose-built systems, like those used in HR, the risks differ slightly according to the intended function. What follows is an analysis of contemporary uses of AI to help fulfil key HR functions, including recruitment, setting compensation, scheduling work, and performance management. Each of these discrete and purpose-built systems for HR functions is assessed in terms of the overall objective that is specified, as well as the data and programming involved.

## Recruitment

Of all human resource functions, recruitment has been most transformed by digitalization, and more recently by AI. The shift to online recruitment began in the mid-1990s and was hailed by industry experts and economists as a revolution that would improve labour market matching by lowering the friction costs associated with job search and worker sourcing (Krueger 2000). Yet the ease of advertising and applying for jobs has increased the volume of applications, making the

task of sorting candidates more cumbersome, thus necessitating technological solutions to assist the process. Recruitment is thus a prime example of the “paradox of automation” (Gray and Suri, 2019), whereby each problem that technology tries to solve, creates a new problem to be solved.

The use of AI (or sometimes just algorithmic systems) in recruitment spans the different stages of the recruitment process, including sourcing of applications, but also their screening and eventual selection. In each of these stages, there are concerns around some or all three parameters of AI systems: objective, data and programming.

## Sourcing

Sourcing refers to how candidates are recruited, including which platforms – job portals and social media sites – jobs are advertised on, and most importantly, who is shown the ad. Employers develop job advertisements and decide which third-party platform to use for distribution. The platform will then display the advertisement to the platform’s users depending on criteria set out by the employer and on how the algorithm is programmed. In general, AI systems are designed with the objective of specifically targeting individuals who are more likely to submit a job application after seeing the advertisement, or potential applicants with specific skills or traits desired by the hiring organization.

Each third-party platform that displays job advertisements has embedded AI systems that dictate to whom the advertisement is shown. The targeting of job advertisements to source candidates is complicated both by what data are available, what data are considered relevant to the objective, and how such data are processed. Platforms have a tremendous amount of data on their users, including demographic information about users’ age, sex, race, sexual orientation, religion, national origin, or political affiliation – the use of which can lead to explicit or implicit forms of discrimination.

The most blatant form of discrimination is when advertisements deliberately target or exclude potential applicants based on these characteristics. In 2002, Meta (Facebook) agreed to a settlement with the United States Department of Justice in which it would eliminate features in its advertising that allowed landlords, employers and credit agencies to discriminate against protected groups. The settlement followed revelations from investigative journalists that Facebook allowed housing marketers to exclude African Americans and others from seeing some of their advertisements (Tobin 2017; Kofman 2022). It is not known whether such features still exist in other countries.

Discrimination can also be embedded in the programming code and how algorithmic optimizations are realized. In other words, an algorithm’s inferences about users’ preferences can also be statistically discriminatory. An experiment undertaken by researchers on Facebook examining how a STEM job was advertised found that although the ad was designed to be gender neutral, fewer women were shown the advertisement (Lambrecht and Tucker 2019). The authors attribute their findings to how the AI system optimized viewership of the advertisement. As they explain, since ad cost minimization is embedded in the code – to get the highest return for the least amount of money – the platform’s algorithm interpreted the code design and ‘learned’ how to more efficiently distribute and display the job advertisement. In this case, the algorithm ‘learned’ that the advertisement would receive more hits per dollar spent if by maximizing exposure to men, as advertisements targeting men are less expensive than those targeted at women; the authors thus deduced that the artificial intelligence evolved in this biased direction. Because programming code is not readily available to the public, however, it is also possible that men were pre-programmed initially as a preferred demographic.

In many jurisdictions and whether intended or not, actions that result in discriminatory outcomes are illegal. While steps have been taken to mitigate this risk, discrimination nonetheless persists. A 2020 audit of Facebook's Ad Library API to ascertain ad distribution published after the Facebook settlement and changes to the advertising portal (Kingsley et al. 2020), found that advertisements for economic opportunities were still not distributed proportionally or equally by gender and age. Women were more likely to receive advertisements for "secretary" or "nurse" while men were more likely to receive ads for construction and STEM jobs. The authors advance three possible explanations. The first is that advertisers evaded the anti-targeting provisions by not disclosing that their ads were housing, employment or credit-related (while targeting ads for these services is illegal, it is permitted to target advertisements for consumer marketing purposes); in this case, advertisers were still able to overtly include gender and age-related data and discriminate on this basis. A second possible explanation is that the ad-distribution algorithm improperly optimized who should see advertisements, leading to systematic biases. In this case, the programming of the AI system 'learned' about the gendered division of labour in certain jobs from how viewers interacted with the advertisements – and then – to maximize efficiencies, preemptively targeted certain demographics for certain jobs. A third explanation offered is that advertisers used proxy variables to target an audience, and these were discriminatory. In other words, discrimination occurred because organizations advertised to viewers based on another characteristic that was so highly correlated with gender or age that they may as well have been targeting viewers based on these protected characteristics (Kingsley et al. 2020).

## Screening, Interviewing and Selection

The screening process typically begins with an algorithmic assessment of a candidate's suitability. Some of the data – such as information provided on applicant CVs – is provided directly by the applicant. Other types of data – like screening tests, which can include a cognitive assessment, a personality test, or a video interview – may be collected by the algorithmic system as part of the application process. Still other sources of data – like an applicant's credit history or social media activity, can be sourced from third parties.

The screening of CVs is done according to pre-defined criteria selected through pattern recognition and the analysis of prior datasets, in addition to criteria determined by the employer. Based on the organization's needs, variables are created to capture the criteria and are then programmed into a software tool that is used to assess candidates' profiles. Applicants are filtered accordingly. While the intent is to eliminate unsuitable candidates, potentially suitable candidates may be excluded if the pre-defined criteria are too narrow in scope or the data used to capture these criteria are poorly conceived or reductive. Ajuna (2023), for example, discusses the case of a firm seeking to hire an engineer. Although the firm would have considered any qualified engineer, the selection algorithm excluded the applicant because he had a Bachelor of Arts degree in engineering, as opposed to a Bachelor of Science degree in engineering, despite graduating from a prestigious university. While this is a low-level risk that can be easily fixed, it illustrates the problems of overly reductive programming and the need for human intervention in the design and implementation of AI systems.

The shift towards greater quantification of recruitment has also led to an expansion in the scope of the data collected and analysed about applicants. While background checks are common within traditional recruitment, employers now have access to an expanded pool of data detailing a worker's background, including information from credit reports and social media activity (Kaur et al. 2020; Christl 2023a). In addition to the privacy concerns from screening workers based on their social media profiles, there are also potential legal risks, if for example a potential candidate is excluded based on past union activity (Davison et al. 2012). There is also the risk that as the



use of social media data becomes more readily used for screening applicants, those applicants who do not have a social media presence will be less likely to be considered.

Once candidates move past the initial screening, they may then face other assessments, including cognitive, personality or situational judgement tests before being interviewed. Prospective hires are asked to engage directly with a digital system, with the assessments sometimes presented in a gamified format (Bodie et al. 2016). Gamified personality tests provide AI systems with vast amounts of data, including video and audio content, response times, answer selection and more. AI systems use machine learning to process this data to predict possible future behaviours. For example, a game requiring job applicants to inflate a digital balloon with air in exchange for a reward was used as a means for identifying propensity for risk-taking behavior and reaction time (Raghavan et al. 2020).

In traditional employment assessments, there is some theoretical reason for associating certain traits with desired characteristics needed for the job, which can be defended if needed (Raghavan et al. 2020). With machine learning, however, the associations that are made between datapoints cannot necessarily be explained, and thus the relationships between the data inputs and conclusions drawn from their analysis are often unclear. Indeed, the lack of theory is an underlying feature of big data analytics (Elragal and Klischewski 2017). In these cases, the fundamental objective of the AI system – such as selecting a candidate who is likely to achieve success in the workplace – has been operationalized by reducing a complicated array of applicant qualities and experiences to a few datapoints. But is this approach reliable, or even valid? How well, for example, can a balloon game approximate the reaction time a worker might have in a real-life situation? And is reaction time an accurate and stable indicator for employee retention? As (Raghavan et al. 2020) argue, the algorithm may be inferring relationships that may not be decipherable and may not hold.

The next stage in recruitment is typically the candidate interview. In AI-led video interviews, interviews are conducted and recorded by the digital system and potential candidates are asked a series of predefined questions by a bot instead of a human (Jaser and Petrakaki 2023). The data collected may include visual data (facial expression, eye movement, hand movement), verbal data (vocabulary, key words), and vocal data (voice tone, pronunciation). These data are then incorporated into algorithmic scoring systems that assign applicants scores and predict whether the candidate is a good fit. At this stage, AI can either assist the interview, by making recommendations that humans can review (AI-assist), or lead, the process where the AI system selects candidates without human review (AI-led). In either case, it is not always evident that the data being collected (e.g., eye movements) captures the objective initially laid out in the recruitment. While it is true that human-led recruitment is also plagued by the problem of whether the information (data) gleaned by the human interviewer corresponds with success on the job, it is also true that replacing the human with the AI system, does not solve the problem.

To understand the reductive nature of some recruitment systems, consider the case of the Treasury Board of Canada, which in 2024 undertook a pilot study to use AI in its hiring practices. The Board presented the technology as a means to increase fairness, reduce bias, improve speed, and improve the candidate and job fit. As a public institution, the Treasury Board completed an Algorithmic Impact Assessment to identify possible risks from using the system. The company that they partnered with, a startup based in Toronto, Ontario, filed a patent on their algorithmic evaluation mechanism which consisted of two components: a behavioural assessment classification and a ranking algorithm. The patent documents provide significant insights into how the system functions.

The client had identified as its objective the selection of candidates with a “growth-mindset”. As a result, the AI system was programmed to identify and score candidates according to whether the vocabulary used in their initial CV and during the interview corresponded to persons with such characteristics. During a recorded interview, candidates were presented with targeted questions that had been developed to elicit answers featuring specific words identified as being associated with a growth mindset. Respondent answers were transcribed using voice-to-text software.

To analyze the transcribed interviews and generate a behavioural ranking, the system was programmed to look for the target words or phrases in the candidate’s response that matched the predetermined word list. When evaluating a candidate’s “growth-mindset”, the system would generate a sum behavior probability based on whether the words were used. More frequent use of words like ‘growth’, ‘development’, and ‘learning’ determined a greater “quantity” of the “growth-mindset” behaviour. The system allowed for calculating a quantity score of the behaviour at the level of the sentence, the question, or across the entire transcript – which is to say – for the candidate-as-a-whole. These results were then used as an input to inform decision-making in the hiring process.

The approach raises a fundamental concern – true of human-led interviews as well – about the supposition that responding with answers that suggests a “growth mindset” embodies the traits needed to be a successful employee. This fundamental weakness in recruitment has not been eliminated by using AI. Indeed, another case study of the experience of a large multinational in integrating AI into its recruitment practices highlighted “discrepancies between the envisioned machine learning (ML) approach and the complexity of knowledge work, such as the inability of the system to know what defines a good employee” (van den Broek, Sergeeva, and Huysman 2021, p.1574). For this company, it took two years to develop the ML system as the first pilot revealed divergences between the programme’s identification of promising candidates and managers’ assessments of the candidates. The developers realized that they needed to better understand managers’ reasoning in the hiring process, thus hiring managers were requested to document their reasons for choosing or rejecting candidates, and with this data, the programme was re-trained. But even then, there were shortcomings, eventually leading to a hybrid approach, and one that used linear techniques, rather than neural networks to ensure explainability (van den Broeck et al., 2022).

## Compensation

Determining and implementing a workforce compensation system is a key human resource function. Organizations are challenged to develop and implement a compensation strategy that will attract and retain a high caliber workforce while meeting the organization’s goals (Gomez-Mejia and Welbourne 1988). Such a strategy can also aid in performance management; by linking a worker’s compensation to individual or firm performance, organizations can reward behaviors that align with their overall goals and improve efficiency (Yanadori and Marler 2006).

When establishing a compensation system, organizations are tasked with considering a wide range of pay features. This includes choosing a mechanism for structuring worker pay – such as adopting a salary or hourly pay system, or incentive-based systems such as those determined by commissions or performance. Depending on the country’s labour laws, employers may also need to assess what benefits will be included in compensation – such as pension or retirement provisions, paid time off, or healthcare. In addition to compensation structure, employers must also select how they will determine pay levels. This exercise often considers how jobs are valued internally – or relative to others within the organization; externally – within the market more



generally; and individually – to reflect the skills and performance of an individual worker relative to others (Marler 2024). While the use of AI systems within the field of compensation has not yet become widespread, systems are being developed and adopted that are influencing how compensation is determined.<sup>1</sup>

That compensation is a target area is perhaps no surprise given that it accounts for such a large percentage of total economic costs (Gerhart and Rynes 2003), and the information is quantifiable. Firms wishing to maximize efficiencies in this area will likely adopt an AI objective with aims like reducing compensation costs, attracting or retaining the desired workforce, and improving motivation or rewarding productivity.

To date, companies that have promoted the use of AI for determining worker compensation have been those with access to large volumes of pay-related data. This has included large companies that generate significant amounts of compensation data internally, and – more frequently – software companies offering compensation services and third-party payroll providers who aggregate data across organizations (Marler 2024). In the United States, but also other countries, payroll provision is concentrated among a few companies. For example, Automatic Data Processing Inc., which owns a host of HR-related software packages under its ‘Workforce Now’ software suite, is a central processor of payroll data in North America, processing payroll data for one out of every six workers in the United States.<sup>2</sup>

There are a growing number of AI-informed compensation software tools on the market that have been developed to complement the hiring process, for example, by forecasting salary projections of job candidates based on an evaluation of their resumes and skills (Gong et al. 2022). Built with data from across the job market, these AI systems can generate a prediction of the cost of compensating the candidate now and in the future. In doing so, they provide the hiring organization with information about the expected cost of attracting and retaining the individual that can be used as a benchmark for negotiation, if a job offer is made. Sometimes these tools are accessible to organizations that have outsourced payroll and compensation tasks. In some instances, organizations are, “willing to make their [compensation] data available [...] in return for access to the predictive models and benchmarked comparisons” generated by third-party companies that have analyzed or processed data from a wide variety of firms (Tambe et al. 2019).

Providers of these types of AI programmes claim that benchmarked comparisons can help firms to determine internal wage grades and wage ranges that are consistent and competitive with market norms. While wage and compensation data are significantly more objective than variables aiming to capture more complicated phenomena like performance or predicted performance, issues with AI programming still arise in the field of compensation. In this case, “individual employers [must determine whether] their context is distinct enough that an algorithm built on data from elsewhere will make effective predictions in their own organization” (ibid). Put differently, if a software system recommends or issues an average compensation package for a group of workers based on an analysis of industry norms, this recommendation is only useful if the organization using the system and the workers for whom the compensation package has been developed are representative of industry norms.

<sup>1</sup> More common has been the use of AI for benefit administration. Many companies use AI, for example, to automate reminders that are sent to workers about enrollment in benefit programmes. Or AI is used to analyze individual and family demographics to suggest the best benefits plan for a specific worker. Predictive systems can also be used to help employers to plan and maximize benefit use; they can also look for patterns and trends, identifying errors or potential instances of fraud. In this instance, AI systems are automating HR functions previously fulfilled by HR professionals.

<sup>2</sup> [www.adp.com](http://www.adp.com).

The inferences and patterns drawn from large datasets can also be used to craft individual compensation packages. For example, HRbrain.ai, a company developing compensation software promises to “personalize incentive plans and bonuses using insights on what motivates each employee” (HRBrain.ai 2024). However, these types of predictive inferences rely on imperfect measures and are limited in scope, often failing to capture complex phenomena like performance, or individual preferences and motivations. Not all workers, for example, will be motivated by the same pay structures. For example, some workers paid hourly may be motivated by premium wage rates for working overtime or, if salaried, may be incentivized to work from home or on weekends to receive a bonus that is linked to achieving specific work targets or outputs. For each worker, the precise level of compensation required to incentivize performance may differ. Workers with debt or large financial commitments may readily take any available work because they are motivated by compensation, while others who have non-work time commitments or interests may be less financially motivated and might be more inclined to accept a compensation package that has a higher number of vacation days. If the relevant variables for workers whose compensation is determined with the assistance of AI are not incorporated into the program, the outputs generated by the system may not reflect the workers’ true priorities or – consequently – the most efficient deployment of resources for the firm.

Among the most robust uses of AI for compensation has been in the gig economy, where large, on-demand firms have access to huge volumes of data that are structured on an individual basis and at a granular level. In managing and compensating a workforce comprised of individuals classified as independent contractors, several large firms have established performance-based variable pay under an employment regime that is beyond the purview of minimum wage laws in most jurisdictions (Dubal 2023). Under these systems, total pay is a function of access to work, pay rates, and incentive structures – all of which rely on individual worker data and thresholds for acceptance (Dubal 2023; Teachout 2023). The use of data-driven pay setting in gig work is important, given the possible expansion of this practice into traditional industries.

When using AI to determine compensation, the larger risks stem from programming. For companies, precisely tailoring wage rates to individual workers in real time according to worker preference and current market conditions can help firms to capture a greater profit share; yet such systems are also discussed as a type of ‘algorithmic wage discrimination’ (Dubal 2023). Profit-seeking firms are motivated to reduce costs where possible, and for lean and asset-poor gig economy firms, labour costs present the largest opportunity to do this. Unlike traditional compensation and minimum wage legislation that has sought to create a floor to provide all workers with economic security, algorithmic wage discrimination functions in reverse. The system begins by offering workers a range of job opportunities; then, over time, as workers accept or reject jobs based on a range of data points such as amount, length of time the job will require, and other individual considerations that reflect individual workers’ needs and preferences, the AI system learns who will work for what amount of money and under what conditions. The concept refers to algorithmic wage *discrimination* because once the AI system has identified workers willing to work for less, the system has been purported to limit the earning potential of that very group (Dubal 2023).

A key risk of personalized pay systems or using AI for compensation more broadly is that pay rates determined with the assistance of AI may run afoul of pay equity provisions. If flawed models and programming are used to establish pay rates and structures, for example, employers may be liable even if they did not develop the system (Markel et al. 2023). AI systems that generate highly individualized compensation recommendations may be particularly prone to equity concerns because their introduction often represents a shift away from more centralized and often standardized compensation regimes. Centralized governance of pay has aims to ensure that all

workers will receive equal pay irrespective of their demographic characteristics – helping to overcome systemic inequalities that have not only historically paid select worker less – but have also helped to shape cultural norms and expectations among those workers about how much their labour is worth, and thus, how much they will accept. Extensive research on the gender disparity of wages, for example has shown that women not only earn less, they are also more likely to accept lower wage offers compared to men with equal credentials (Recalde and Vesterlund 2022). Within the context of individualized compensation, machine learning algorithms could translate the tendency for women to have lower earnings expectations and to be paid less overall, for example, into lower wage offers on a systemic level.

The opacity of AI system development and use for compensation is also a particular concern given that pay transparency has long been recognized as necessary for promoting pay equity (Chicha 2006). In recent years, pay transparency legislation has been introduced in about half of OECD countries to mitigate bias embedded in hiring and pay determination, and thus combat the gender pay gap. While pay transparency has not been a panacea, there is evidence that it has led to a narrowing of the gender pay gap (Baker et al. 2023). Among other things, transparency regulations have helped women to better understand salary structures within their workplace so that they can better advocate for themselves or to raise grievances or challenges when they feel discrimination has occurred (Duchini et al. 2024). By shedding light publicly on organizations' compensation packages, pay transparency rules push firms to proactively address pay discrepancies where they exist.

When it comes to AI, some developers claim that AI's ability to identify compensation patterns and trends can help strengthen pay equity by identifying biases introduced by imperfect human-driven compensation systems. However, when AI used for compensation purposes lacks transparency and relies on proprietary data and code that are beyond the purview of oversight, the systems may have the opposite effect and undermine the recent progress that has been made in pay equity.

## Scheduling work

Perhaps the oldest application of algorithms to HR functions has been related to the “science of ‘scheduling’” (Weaver 2006). First introduced in the 1950s via a concept called, ‘Critical Path Analysis’ (CPA), researchers advocated for the adoption of a quantitative approach to business decision making that could help to analyze, plan and schedule large, complex projects. As explained in a 1963 edition of the *Harvard Business Review*, this approach “provides a means of determining (1) which jobs or activities, of the many that comprise a project, are ‘critical’ in their effect on total project time, and (2) how best to schedule all jobs in the project in order to meet a target date at minimum cost” (Levy et al. 1963). In short, efficient utilization of resources was the historic goal of – and answer to – the ‘scheduling problem’.

Since CPA was first introduced various iterations of context-specific ‘scheduling problems’ have been examined. In each instance, the objective remains constant: to optimize the use of resources given a predetermined set of constraints that reflect both the nature of the ‘job’ and the conditions under which the job is performed (Graham 1966; Jain and Meeran 1999). Constraints are highly context dependent. For example, in simple manufacturing or assembly situations, the only constraint may be workflow order; in more complex configurations, task completion is restricted to a worker with a certain skillset; firms must juggle competing tasks with different time constraints; or a single machine may present a bottleneck for production. When operationalized, these constraints are transformed into variables that are then incorporated into the AI's

software code. There are no limits to the potential number of constraints shaping scheduling problems; however, in each iteration, there are a limited number of potential solutions, and the goal remains to identify the best and most efficient option. The solutions, in turn, shape the daily experiences of workers – including the time and duration of their shifts, and the distribution of tasks among the workforce.

Since scheduling problems became a field of interest and inquiry, their mathematical precision has rendered computers and programming useful for finding solutions. Indeed, the 1959 paper that first re-envisioned project planning and introduced Critical Path Analysis was motivated by using computer-oriented systems to improve efficiency; the paper even mentions algorithms as part of the method for determining solutions to these types of problems (Kelley Jr and Walker 1959). When this paper was first published, the authors anticipated that fully implementing CPA would require “large-scale computing equipment”. As computational power has increased, machines have been able to accommodate a larger volume of constraints of scheduling problems and – in the AI era – are now able to undertake analysis while accommodating real-time inputs and generate dynamic solutions seemingly instantaneously. It is thus not surprising that since the early 2010s, the use of AI for solving scheduling problems has grown (Del Gallo et al. 2023).

Early AI and algorithmic-scheduling systems are built on probabilistic models reflecting the constraints and goals of the system. Within manufacturing, AI scheduling has been used to reduce production costs, to optimize job completion times, or reduce the time required of HR or other administrative personnel to undertake the work of determining staff and production scheduling (Del Gallo et al. 2023). Within the service sector, the most basic AI and algorithmic scheduling systems largely base predictions on past staffing needs (Chapados et al. 2014). Staffing needs often increase during times of heightened demand. While some fluctuation in demand is predictable and can be anticipated (for example, the bulk of grocery store purchases are completed outside of traditional working hours, and flight demand increases around holidays), demand can also be stochastic (for example, rain may increase the demand for taxi services or public transit). More advanced ‘dynamic scheduling’ systems can be programmed to consider a wide variety of exogenous events impacting organizational needs or performance, and can adapt in real-time accordingly – although no system can foresee all possible challenges (Akridge et al. 2024).

Dynamic systems have also been successfully used across organizations. For example, during the height of the Covid-19 pandemic, the State of Massachusetts in the United States created an algorithm and job portal to re-allocate nursing staff in long-term elder care environments to facilities facing labour shortages. The algorithm considered the urgency, vacancies, and locations of nursing homes; the professional data, location, and availability of nursing workers; and criteria and matching policies that were set by the state. Overall, the use of the system not only helped match workers with jobs, but it also increased the staff to resident ratio and yielded a net increase in the hours worked per day. Those involved with the project noted that the success of this type of initiative, required “close collaboration of nursing homes, policymakers and the centralized matching process development team” (Zarei et al. 2023). Ironically, no mention was made of the nurses affected by the system.

In addition to recommending changes to the total number of workers engaged at a given time by a particular organization or deploying workers across organizations, AI scheduling systems can also shift how labour is deployed within an organization by influencing how tasks are distributed between workers and how tasks are ordered. Achieving this type of dynamism requires granular data about the job process and extensive programming to accommodate a wide range of constraints. Data used in AI systems are now derived from an increased number of sources. These include ‘smart’ machines and tools that are deployed as part of the work process, and data

from the workers. Workers may generate data by working with internet-connected devices or they may voluntarily contribute data in the form of their working time availability or scheduling preferences. In decentralized systems where scheduling inputs come from a range of actors, AI can help ensure that the organization is still able to meet its operational needs (Destouet et al. 2023). However, if system designers overlook crucial aspects of resource deployment or organizational function, or are missing data, there may be significant efficiency losses which undermines the very justification for introducing these systems in the first place.

When it comes to the deployment of labour within the organization, real-time tracking of job needs and progress allows organizations to assign workers to the most urgent tasks and to track their progress and completion. For example, real-time data gathered from airplanes is used to facilitate ground operations and to direct baggage handlers to their next task. The instantaneous data transmitted can update airport personnel in the event of a delayed departure or arrival, and the tasks allocated to airport workers can be updated accordingly (Amiri and Kusakci 2024). AI task allocation is also used within the food service sector, where systems are programmed with “predefined target times for the preparation of each type of menu item” to aid with workload balancing between kitchen team members (Christl 2023a), and in the hospitality sector where systems keep track of demand and inventory, centralize hotel-wide communications and allocate employee tasks, including that of the cleaning staff (Spektor et al. 2023).

But systems can also be designed to support worker autonomy. For example, within the UK, BT Group, a provider of telecommunications services uses AI-powered task allocation systems to coordinate the work of field service agents – a fleet of often mobile worker technicians tasked with system repair and maintenance operations (Wang et al. 2022). The organization shifted from a top-down task allocation previously overseen by managers, to an AI-driven work recommendation process where pre-qualified and trained field service agents were provided a list of tasks in line with the company’s productivity expectations. From this list, technicians could select their job functions and locations. This AI-recommendation algorithm was built in-house by BT Group with input and buy-in from field service technicians and other company stakeholders. With algorithmic explainability at the center of the company’s deployment strategy, management and technicians alike understood that the workforce using the system needed to retain decision-making latitude to deal with unforeseen customer issues and unusual workplace demands that may fall outside of parameters of the code underwriting the AI programming (ibid).

Despite the benefits of building collaborative and transparent systems, there are also a growing array of more advanced and opaque systems used to distribute tasks based on a worker’s predicted performance. In call centres, for example, AI systems are used to create profiles of workers and clients using historical call data; then, by learning from past interactions between these actors and their success or failure, the AI system will match phone partners to achieve whatever the organization’s goal is (sales, assistance, etc.) (Christl 2023b; Kresge 2020). Unlike AI systems that distribute work based on formal qualifications and standardized credentials, however, the kinds of AI systems used in call centres and other service-oriented jobs with a high level of customer-employee interaction, create models of users that are built on personality traits, demographic information, or other data for the purpose of social matching (Terveen and McDonald 2005). These types of jobs require the performance of emotional labour. While companies may establish norms around social interactions, how workers “deliver ‘sincere performances’ and manage emotions according to management prescriptions do not result in a uniform response on the part of employees” (Seymour and Sandiford 2005). Within AI social matching systems, these variations in informal human responses are exploited for the purpose of trying to improve task allocation. However, the systems are relying on data and data inferences that are not scientifically



supported, rendering the connections drawn either unclear or inaccurate.<sup>3</sup> Even for personality tests, commonly used as datapoints in social matching algorithms – but also in hiring, team composition, and other work functions – research on their reliability is conflicting and even discredited (Greenberg, 2024). Nevertheless, personality tests remain readily used in industry and their results can give employers a false sense of confidence by seeming to ground decisions in evidence that is not actually present, in addition to introducing other risks related to privacy and discrimination (Stabile 2001).

Beyond the concerns about the accuracy of predictive allocation – particularly when it is based on personality inferences – there are other significant and serious downsides to collecting the kinds of data these systems require. Many of the drawbacks of real-time tracking are related to the close connection that this type of monitoring has with performance management systems – discussed in the next section.

## Performance management

Among the most extensively researched areas of AI at work has been the use of AI for performance management. Strongly tied to the legacy of Taylorism and the more general principle of scientific management (Birnbaum and Somers 2023), AI performance management relies on close measurement of worker activity and the inferences that can be derived from patterns and trends. Common applications of AI performance management can be used to determine whether workers are meeting output goals, and increasingly are also being used to assess performance quality (ibid).

Within performance management, AI systems are often designed with the objective of improving efficiency and productivity – such as producing more goods or serving more customers with fewer resources. Such goals can be established at a firm level (such as generalized output of the workforce writ large), or, when information on the production process involves granular data about individual tasks, at an individual level (Kellogg et al. 2020). Because of AI's capacity to incorporate and analyze new sources of data in real-time, performance targets – and consequently, the framework against which workers are evaluated – can also change on an ongoing basis (Kresge 2020).

Some occupations and tasks lend well to objective measurement. Jobs that are highly standardized and repetitive are easier to quantify with data-driven performance management techniques. In occupations where workers' decision-making opportunities are constrained, such as on factory floors in industrial manufacturing, worker performance can be measured by units of output. Facilitated by new technologies that quantify work (Bernstein 2018), outputs previously measured by a manual count of units produced can now be counted faster using an array of smart tools and scanners. This creates more opportunities for task and workforce oversight and control, particularly in environments where jobs have already been disassembled into an assembly line comprised of unique tasks (Littler 1978). The system is programmed such that the data amassed from these 'smart' devices are used to monitor the contributions of individual workers. For these standardized and repetitive jobs, the three pillars of objective, data and programming are met, but typically at a cost for the workers who often face high and ever-increasing performance targets.

Among the best-known examples of the application of AI technologies are in the logistics sector. Workers at warehouse fulfillment centres in Italy (Delfanti 2021), North America (Vallas et al. 2022), and likely elsewhere, are subject to a range of tools that quantify their speed and movement. AI systems direct ‘pickers’ where to go in the warehouse to collect items that will be shipped to customers. Workers scan each item after locating it, providing the AI system with the necessary information to demonstrate that a worker has met the intended goal. These systems can then be used to determine average working speeds and to set performance goals for the workforce as a whole or for individual workers (Krzywdzinski and Gerber 2019).

In highly controlled workplaces, the notion of measuring productivity as a function of output remains relatively unchanged from the productivity measures that have existed for decades (see, e.g. (Hiba 1998)). Consequently, the data collected and the analysis derived for these AI systems about worker performance remain mostly unchanged, and the relatively objective measures of performance vis-à-vis these measures continue to be used to reward workers or to discipline them – albeit more quickly and at a more granular level (Kellogg, Valentine, and Christin 2020). Under these systems if a worker is unable to achieve a minimum threshold it may trigger a notification for the worker to speed up or even result in disciplinary action being taken against them (Ball 2022).

While these systems may appear relatively objective regarding how output is measured, the designers of these AI systems make deliberate programming decisions that delineate between productive and non-productive activities. For example, in research conducted in slaughterhouses on assembly-line meat processors, workers were monitored using wearable devices and sensors capable of measuring activities like acceleration, orientation, and location and able to recognize human activities such as walking, standing, sitting, and sleeping. Workers’ activities were then identified according to pattern recognition and classified as either productive or non-productive, with the express goal of reducing downtime and increasing profits (Forkan et al. 2019). The system, however, was not programmed to account for unexpected problems in the assembly line. Moreover, the overall objective comes at a significant cost, as increases in line speed and pace of work are directly associated with increases in occupational injuries and illness.<sup>4</sup>

New sources of data and efforts to develop more complex productivity and performance measures have also emerged. This has created opportunities for the quantification of work and data-driven management to spread into new fields. For example, research on the trucking industry has found that despite lower levels of routinization within this job (employees face new geographic locations daily), high levels of tracking and surveillance have now become industry standards. The data collected are vast and typically include “a driver’s fuel efficiency and idling time, speed, geolocation and geofencing (notifying a dispatcher if a truck has departed from a predetermined route or arrived at a terminal), lane departures and braking/acceleration patterns, cargo status (e.g., the temperature of a refrigerated trailer), and vehicle maintenance/diagnostic information” (Levy 2015). AI systems use these data with expectations about industry standards or performance to develop new productivity measures for a wider range of items beyond merely distance driven.

Traditionally, occupations that require complex problem solving or are characterized by high levels of interpersonal interaction have not been seen as well-suited to the types of quantifiable performance metrics detailed above. However, recent developments in natural language processing technologies and facial recognition have pushed AI development towards “emotional

<sup>4</sup> At over 20%, injury and illness rates in the meatpacking industry are among the highest in North America, and have been attributed to increases in line speed and the pace of work (Cai et al. 2005; Kyeremateng-Amoah et al. 2014).

recognition”. The data collected on interpersonal interactions are used to offer workers real-time feedback on how they are doing their jobs. Voice recognition used on call centre workers, for example, will ‘listen’ to phone calls and generate recommendations for workers regarding the call content and their emotional disposition (Christl 2023b). Such recommendations are driven by what the AI system predicts will result in a favorable call outcome. The higher level of workplace control has potentially negative implications for job quality. Research on job quality has found that when workers are able to exercise decision making latitude, like having control over the order of tasks, the speed at which they work, or having an ability to influence how the task is associated with more favorable perceptions, improves job quality (Green 2007). Performance management AI used in call centres strips workers of this latitude by not only dictating what workers must say, but also by closely monitoring the number and length of calls that they complete, and the relative outcomes of each (Kresge 2023, 20). Like the issues outlined in the section above, these systems may rely on or identify variables that are not clearly related to workers’ performance, making the information gleaned difficult to interpret.

The growth of AI-driven performance indicators has exploded since the Covid-19 pandemic. With the shift to remote work, many knowledge workers whose work performance was previously not monitored or measured found themselves working from home. Unable to see the workforce, employers began to rely on a host of monitoring tools as a stopgap measure to assuage their concerns about shirking. The systems collect data on keystrokes, mouse movement, screen activity, email content and more (Aloisi and De Stefano 2022; Ball 2021). Major software providers like Microsoft have included productivity metrics into their software applications. By logging workers’ activities on programs like Microsoft Teams, programs will track not only how much time a worker spends online but also will make inferences about the quality of their engagement. Functionally, this method of operationalizing performance evaluation for AI systems uses variables about *what* workers do to draw conclusions about *how well* they do it. On the application Microsoft Teams, for example, the company provides analytics about how employees use the platform. This is marketed as “Performance Analytics” in Microsoft’s marketing materials. Administrators can access details such as whether a worker is actively in front of their computer, away, or busy – or to track whether employees engage in conversations or file sharing. Yet an overreliance on this data can leave employers with a poor understanding of what is really happening in the workplace that is not coherent with the stated objective of the system. A worker, for example, can be away from the computer and still be engaged in productive activities; and the relationship between two workers with a high level of engagement could be collaborative or it could be combative. In this case, the problem with the use of these AI systems is that users draw conclusions about performance based on data that are not necessarily performance related, and variables that do not actually capture performance quality.

As a result, some of the performance metrics provided may be difficult to interpret both for managers and workers. Rahman (2021) documents the experience of freelancers on a well-known freelancing platform who were affected by a radically revised scoring algorithm, designed to address ‘rating inflation’ on the platform. The new system was introduced unilaterally and without notice and did not explain how scores were calculated, as the platform did not want freelancers to game the system. This meant, however, that the workers had no means of knowing what they were doing well or not, or how to improve their performance (Rahman 2021).

But even in traditional workplaces, where there is possibility for feedback from managers that can support professional development, managers might find that they are not able to guide staff based on the metrics collected, especially as the systems may lack transparency regarding the variables used and what data are collected and monitored. As such, the information is



either of no use or worse, risks steering performance in an ineffective way and weakening the role of managers.

### ► 3 AI in HRM: Unbridled optimism

---

For HR to fully reckon with the opportunities and risks of AI, the field must reckon with an underlying positivist orientation and bias towards numeracy as emblematic of a superior way of knowing and understanding the world of work. These tendencies within the field have yielded an unbridled optimism about the benefits of AI that is largely devoid of an understanding of the three pillars that underpin these systems and their implications. A failure to reckon with the assumptions that have fueled HR's embrace of this technology puts in peril the fundamental mission of HR of effectively hiring, retaining, motivating and managing workers to secure organizational success.

While some facets of the workplace are arguably more objective and well-suited to quantification – for example, the number of goods produced, or the number of hours that a worker is scheduled for – HR functions are incorporating AI and quantification for a range of other objectives that evaluate workers against complex constructs. When applied to performance management of knowledge workers or candidate predictions in recruitment and selection, human behaviour cannot be broken down in to discrete variables and evaluation of certain performance metrics does not guarantee success (Giermendl et al. 2022). Complex constructs represent areas where AI and data-driven management are more likely to generate risks for organizations, to fail, or to produce overly generalizable or not useful outcomes. These models continue to rely on numerical representations that necessarily, “discount[s..] the complexities of the human condition” (See (Wood and Kispál-Vitai 2021) citing Gold 2001: 192), and are thus poorly positioned to predict the nuances that emerge in human social interaction.

Within positivism, numeracy is attractive because it invokes the idea that there are correct and incorrect answers, thereby advancing the notion that there is a unitary and replicable ‘best practice’ for any given task or challenge. Thus, HR’s positivist paradigm also predisposes the field to adopting AI ‘solutions’ that could introduce bias or could be situationally inappropriate. The notion of a universal truth, which is a tenet of this ontological orientation, may lead HR researchers to believe that the same AI system will work in all contexts. Yet research does not bear this out. When a system is developed for a specific objective, but deployed in different environments, AI systems sometimes generate different outcomes. This can be attributed to the training data that are fed into the system or to the variables that are considered by programmers.

As HR professionals work to minimize liability and risk, they will have to determine if AI use is appropriate and what AI system adaptations may be necessary. Specifically, they will need to ensure that any AI system put in place is based on a measurable and clearly defined objective, uses appropriate data, and is programmed to achieve its stated goal. While this sounds straightforward, in practice, it is full of problems, particularly in managing humans at work. This does not imply that there is no role for AI systems in HR management, but rather that there are important limitations that need to be understood and considered when managers embark on their AI journey.

The first challenge of defining the “objective” seems straightforward – selecting qualified candidates in recruitment, setting compensation at a fair price that reflects market conditions but also individual attributes, improving worker performance, to name a few. In practice, however, most of these objectives cannot be sufficiently well defined, or defined in a way can be distilled into data points and then programmed into algorithms. The challenge can be seen clearly in the case of recruitment (both screening and selecting the right candidate), which has always been a

problematic task (i.e., selecting from a pool the individual who is most likely to be the best employee for the job). There is no doubt that AI can save tremendous amounts of time, and perhaps expense, in screening CVs or even interviewing candidates, but whether doing so achieves the objective of selecting the best candidate is less evident. Indeed, what the analysis has revealed is that while the second parameter of AI systems – data – has often been cited as the Achilles' heel of recruitment, due to concerns over possible discrimination, it is instead the inability to operationalize the objective that is the principal concern and challenge.

Problems with *objective* are also seen in performance management. For some routine, standardized occupations, AI systems often have a clear objective, and the quantification of worker movements does provide a relatively accurate data source. Yet the performance targets may be set at such a high level that they imperil workers' safety or well-being. For other types of work that are non-routine, the objective of "managing performance" cannot be met if only some aspects of the work can be quantified and monitored and the data being collected (screen time, email content) do not accurately capture performance. Moreover, performance management systems may not give the intended results or may generate results that are opaque and that do not provide useful feedback to staff (e.g., Rahman, 2021).

With respect to *data*, much criticism surrounds the representativeness of data. This problem can be solved, some contend, by incorporating more data into training the system from a wider cross section of the population. Yet while including more representative data is a valid endeavour, this 'solution' fails to examine whether the data are suitable for operationalizing the objective. As discussed above regarding AI in recruitment, one firm used a system whereby candidates participated in a game of blowing up a digital balloon as a data point in scoring candidates. One could incorporate more representative training data of people's scores in this game, but that would still not address whether data from a digital balloon game constitutes a predictable indicator of performance, and thus whether candidates should be required to generate this data in the first place. In scheduling work and performance management, the only data that can be collected and used to train the system are data that leave a digital trace – not all the work that is performed.<sup>5</sup> This is a problem for knowledge work, as online activity does not capture the breadth of tasks involved, but it is also true of many types of production work, particularly where there are unexpected situations that arise.

For *programming*, concerns center around how the algorithms, which are the key decision-making feature of AI systems, are defined and whether machine learning processes are opaque in ways that elude their users. In recruitment, the discriminatory risks of these systems are not just a concern for workers who may experience discrimination, but also the employers who deploy such systems. In these instances, firms may not only forego promising candidates during the recruitment process, but also might be exposing themselves to discrimination lawsuits as employers remain responsible for executing HR functions without bias (Adams-Prassl et al. 2023). Because of the proprietary nature of the recruitment software, which is often developed externally to the firms that use it, employers will have a hard time ascertaining or explaining how decisions are made. A similar lack of transparency often emerges in the context of AI machine learning systems and can even affect internally developed systems. In these instances, machine learning AI systems are provided with an objective but are left to determine (often through complicated self-learning and multi-layered opaque programming) what variables are 'relevant' and what relationships exist between the two. When this approach replaces hypothesis-driven inquiry

<sup>5</sup> Many functions that workers perform in their jobs are not explicit, just like many of the skills that workers bring to a job involve tacit knowledge (Suchman 1995). It is for this reason that one of the most disruptive forms of labour protest are "work-to-rule" campaigns, in which workers perform only what is in their job description and nothing more (Levitsky and Ziblatt, 2018).

informed by prior knowledge, research or theory, the results can be difficult – if not impossible – to interpret and correlation between otherwise unconnected features may be mistaken for causation. Among the many risks this opacity presents is in the area of pay equity and compensation where a growing number of jurisdictions have introduced pay transparency legislation (Benedí Lahuerta 2022).

While these issues can emerge in any AI system, third-party software can exacerbate them. As AI and software firms compete to develop the most innovative and reliable software systems, their market advantage is often determined by their data sets and programming prowess. Third-party software developers thus maintain a high level of secrecy on how their programmes are trained and coded. While such secrecy helps to protect their interests in the AI product marketplace, it also means that users of these systems have a poor understanding of how they are designed and how they work.

## ► 4 What is an HR manager to do with AI?

---

An empirically driven evaluation of AI systems must originate with a different conceptualization of technology than the narrative that dominates investor and innovator circles. First and foremost, we should reconsider the presumption that AI systems and technology are always useful for the management of people at work. Or that AI systems are ‘easy to use’. These illusions fuel the high-speed adoption of AI in ways that generate risk for employers and workers alike. It is the responsibility of the HR professionals to determine whether, and under what conditions, the use of AI systems is appropriate, and – if AI systems are used – to integrate them into workplaces in a manner that minimizes risks.

Over the past thirty years, research on technology adoption has attempted to identify the conditions under which users adopt new technologies. Within this expansive literature, two features are viewed to be driving forces. These are first, the ease of use of a system – in other words whether the use of a system is free from effort –, and second, the perceived usefulness of a system, in other words, whether the user perceives the system to be useful to aid them in the performance of their job (Davis et al. 2024). The first element has been the focus of user experience and human-computer interaction research; the second – deemed to be the more important factor driving adoption – is largely affective and attitudinal. These are the two foundational principles of the ‘Technology Acceptance Model’; since it was first introduced in 1986, the model has evolved to additionally include a wide range of determinants that drive perceived usefulness (among them, subjective norms, image, job relevance and result demonstrability), determinants for perceived ease of use (e.g. perception of external control, computer playfulness, objective usability), and also moderating conditions including user demographics (Davis et al. 2024).

Research based on the technology acceptance model and its later iterations has shown that we are more likely to adopt technologies when we perceive them to be useful for improving our job performance and when we find them easy to use. This framework has been used by software designers to promote the adoption of software; we should thus assume that the user experience of HR software has been developed with these considerations in mind.

When software engineers consider ease of use, they often focus on user experience and user interface. In many of the AI for HR applications, the technology systems and computer programmes operate with a high level of autonomy and automation that require little of those who deploy them. Indeed, research has found that although only a minority of HR professionals – including those employed in HR analyst roles – feel confident in their ability to conduct predictive HR analytics, many firms nonetheless adopt decisions driven by this type of analysis (McCartney et al. 2021). They achieve this by outsourcing analysis to specialists or – in the case of the growing number of off-the-shelf systems – by procuring ‘easy to use’ prepackaged technology. Herein lies a key problem: AI for HR systems and technologies are easy to use – but many users of these systems have a poor understanding of what, exactly, they are using or how it works.

It is precisely this lack of technological understanding that leads to the use of less-than-optimal systems. Overcoming this shortcoming requires that HR managers participate in the design, implementation and oversight of AI systems. Moreover, for HR functions that pertain to how work is conducted and evaluated (i.e., work scheduling and performance management), it is also necessary to involve the workers themselves. Involving all stakeholders in the design of AI systems provides a means to address all three pillars of AI systems: to clarify the objective,

to ensure that the data used are suitable for that purpose and context as well as representative and transparent, and to ensure that the programming accounts for various conditions that the system is likely to encounter.

Having HR professionals meaningfully engage in the development process will lengthen the development and implementation time and it will also require that HR professionals develop more comprehensive and deeper understanding of AI. This investment is both necessary and worthwhile. A good example of such an engagement was that of the large multinational discussed earlier that spent two years and numerous iterations developing an AI-driven recruitment system, and in the end, adopted a hybrid system (van den Broeck et al., 2022). The company made the necessary investments in time and resources, and fully involved its HR professionals into the process of developing a recruitment software that could provide useful and explainable results. This kind of consultation was key to achieving buy-in from users and cultivating trust in the workplace.

The experience of BT, the UK telecommunications company with respect to work scheduling, is similarly illustrative. The company involved its field technicians in designing a system that combined the benefits of machine learning (knowledge from field technicians' past assignments) while enabling the users, in this case the field technicians, control in organizing their work. The system led to a 10 percent improvement in productivity and improved the wellbeing of the engineers, with mental health absence reduced by over one-third (Wang et al. 2022).

Both of these examples illustrate the importance of moving beyond the superficial "ease of use" advocated by the technology acceptance model to a participatory, multi-stakeholder process to design and build AI systems in which the users, including HR professionals understand not only how the systems work, but how best to use them.

## ► Conclusion

---

According to UN Trade and Development, investments in artificial intelligence are expected to expand twenty-five-fold from 189 billion US dollars in 2023 to 4.8 trillion US dollars by 2033 (UNCTAD, 2025). Human resource management is not immune to this surge in investment. Building on the push towards greater empiricism and new sources of real-time data at the workplace, the incorporation of AI into the different core functions of human resource management has been posited as a natural next step following “people analytics”. Indeed, the field is particularly susceptible to the drive to incorporate AI as it has long sought to improve its standing within the executive suite, favouring empiricism whenever possible. Advocates contend that the predictive capacity of AI is an important advance in HR’s ability to realize its mission of improving and proving the effective allocation of human resources.

But there are limits to how well AI can manage the messy business of managing people at work. Across the myriad of HR functions examined – recruitment, setting compensation, scheduling work and managing performance – there are problems in defining objective, in the data used to train the systems and how it is programmed. These problems are particularly salient when workplace tasks require high levels of human interaction, when work is unpredictable, or when it can be influenced by many variables, or concerns complex human performance and tasks.

Like many areas of the world of work, AI is and will continue to impact HR management. But if these systems are to be useful and achieve the overall goal of effectively managing people within organisations, HR managers must understand how the AI systems work before implementing them. This requisite understanding is particularly difficult when off-the-shelf technologies are used, where the risks related to ill-defined objectives, unsuitable training data, and untransparent programming are likely to loom large; indeed, these off-the-shelf AI technologies may be ‘easy to use’, but they likely won’t deliver satisfactory results. The most effective systems will require the involvement of their users – in this case the HR managers and professionals, as well as workers. By participating in system design, and building and refining the systems, HR professionals can gain better clarity on whether the objectives they are pursuing can be achieved with an AI system, whether the training and real-time data are appropriate, and whether the programming allows meaningful results. As organizations around the world move forward with AI, we hope they will come to recognize the importance of a deliberate, human-centred approach to its use in HRM.

## References

---

- Adams-Prassl, Jeremias, Reuben Binns, and Aislinn Kelly-Lyth. 2023. "Directly Discriminatory Algorithms." *The Modern Law Review* 86 (1): 144–75. <https://doi.org/10.1111/1468-2230.12759>.
- Akridge, Hunter, Bonnie Fan, Alice Xiaodi Tang, Chinara Mehta, Nikolas Martelaro, and Sarah E Fox. 2024. "The Bus Is Nothing without Us: Making Visible the Labor of Bus Operators amid the Ongoing Push Towards Transit Automation." *Proceedings of the CHI Conference on Human Factors in Computing Systems* (New York, NY, USA), CHI '24, May 11, 1–16. <https://doi.org/10.1145/3613904.3642714>.
- Aloisi, Antonio, and Valerio De Stefano. 2022. "Essential Jobs, Remote Work and Digital Surveillance: Addressing the COVID-19 Pandemic Panopticon." *International Labour Review* 161 (2): 289–314.
- Alvesson, Mats. 2008. "Critical Perspectives on Strategic HRM." *The Routledge Companion to Strategic Human Resource Management*, 52–67.
- Amiri, MA, and AO Kusakci. 2024. "A Scoping Review of Artificial Intelligence Applications in Airports." *Comput. Res. Prog. Appl. Sci. Eng. Trans. Ind. Eng* 10: 2900.
- Baker, Michael, Yosh Halberstam, Kory Kroft, Alexandre Mas, and Derek Messacar. 2023. "Pay Transparency and the Gender Gap." *American Economic Journal: Applied Economics* 15 (2): 157–83. <https://doi.org/10.1257/app.20210141>.
- Ball, Kirstie. 2021. "Electronic Monitoring and Surveillance in the Workplace." *European Commission Joint Research Centre*.
- Ball, Kirstie. 2022. "Surveillance in the Workplace: Past, Present, and Future." *Surveillance & Society* 20 (4): 455–61.
- Benedí Lahuerta, Sara. 2022. "EU Pay Transparency Legislation to Address Gender Pay Inequity: What Is on the Horizon and Its Likely Impact in Ireland." SSRN Scholarly Paper No. 4323032. Social Science Research Network, December 21. <https://doi.org/10.2139/ssrn.4323032>.
- Bernstein, Phillip. 2018. *Architecture | Design | Data: Practice Competency in the Era of Computation*. Birkhäuser.
- Birnbaum, Dee, and Mark Somers. 2023. "Past as Prologue: Taylorism, the New Scientific Management and Managing Human Capital." *International Journal of Organizational Analysis* 31 (6): 2610–22.
- Bodie, Matthew T., Miriam A. Cherry, Marcia L. McCormick, and Jintong Tang. 2016. "The Law and Policy of People Analytics." SSRN Scholarly Paper No. 2769980. Social Science Research Network, March 1. <https://papers.ssrn.com/abstract=2769980>.
- Boon, Corine, Deanne N Den Hartog, and David P Lepak. 2019. "A Systematic Review of Human Resource Management Systems and Their Measurement." *Journal of Management* 45 (6): 2498–537.
- Broek, Elmira van den, Anastasia Sergeeva, and Marleen Huysman. 2021. "When the Machine Meets the Expert: An Ethnography of Developing AI for Hiring." *MIS Quarterly* 45 (3): 1557–80. <https://doi.org/10.25300/MISQ/2021/16559>.



Cai, Chunbo, Melissa J Perry, Gary S Sorock, et al. 2005. "Laceration Injuries among Workers at Meat Packing Plants." *American Journal of Industrial Medicine* 47 (5): 403–10.

Cayrat, Charles, and Peter Boxall. 2023. "The Roles of the HR Function: A Systematic Review of Tensions, Continuity and Change." *Human Resource Management Review* 33 (4): 100984.

Chapados, Nicolas, Marc Joliveau, Pierre L'Ecuyer, and Louis-Martin Rousseau. 2014. "Retail Store Scheduling for Profit." *European Journal of Operational Research* 239 (3): 609–24.

Chicha, Marie-Thérèse. 2006. *A Comparative Analysis of Promoting Pay Equity: Models and Impacts*. DECLARATION/WP/49/2006. <https://citeseerx.ist.psu.edu/document?repid=rep1&type=pdf&doi=f376b884cf7bca4f1eec824999a5310f78851f13>.

Christl, Wolfie. 2023a. *Monitoring Work and Automating Task Allocation in Retail and Hospitality*. [https://crackedlabs.org/dl/CrackedLabs\\_Christl\\_Work-Retail-Hospitality.pdf](https://crackedlabs.org/dl/CrackedLabs_Christl_Work-Retail-Hospitality.pdf).

Christl, Wolfie. 2023b. "Surveillance and Algorithmic Control in the Call Centre." *A Case Study on Contact and Service Center Software, Automated Management and Outsourced Work*. *Cracked Labs*.

Davis, Fred D, Andrina Granić, and Nikola Marangunić. 2024. *The Technology Acceptance Model: 30 Years of TAM*. Springer.

Davison, H. Kristl, Catherine C. Maraist, R. H. Hamilton, and Mark N. Bing. 2012. "To Screen or Not to Screen? Using the Internet for Selection Decisions." *Employee Responsibilities and Rights Journal* 24 (1): 1–21. <https://doi.org/10.1007/s10672-011-9178-y>.

Del Gallo, Mateo, Giovanni Mazzuto, Filippo Emanuele Ciarapica, and Maurizio Bevilacqua. 2023. "Artificial Intelligence to Solve Production Scheduling Problems in Real Industrial Settings: Systematic Literature Review." *Electronics* 12 (23): 4732.

Delfanti, Alessandro. 2021. *The Warehouse: Workers and Robots at Amazon*. Pluto Press.

Destouet, Candice, Houda Tlahig, Belgacem Bettayeb, and Bélahcène Mazari. 2023. "Flexible Job Shop Scheduling Problem under Industry 5.0: A Survey on Human Reintegration, Environmental Consideration and Resilience Improvement." *Journal of Manufacturing Systems* 67: 155–73.

Dubal, Veena. 2023. "On Algorithmic Wage Discrimination." *Columbia Law Review* 123 (7): 1929–92.

Duchini, Emma, Stefania Simion, and Arthur Turrell. 2024. "A Review of the Effects of Pay Transparency." In *Oxford Research Encyclopedia of Economics and Finance*. <https://doi.org/10.1093/acrefore/9780190625979.013.860>.

Elragal, Ahmed, and Ralf Klischewski. 2017. "Theory-Driven or Process-Driven Prediction? Epistemological Challenges of Big Data Analytics." *Journal of Big Data* 4 (1): 19. <https://doi.org/10.1186/s40537-017-0079-2>.

Enderes, Kathi, and Matthew Shannon. 2019. *People Analytics Solutions: Market Primer*. Marketplace Landscape. Deloitte.

Fleetwood, Steve, and Anthony Hesketh. 2008. "Theorising Under-theorisation in Research on the HRM-performance Link." *Personnel Review* 37 (2): 126–44.

Forkan, Abdur Rahim Mohammad, Federico Montori, Dimitrios Georgakopoulos, Prem Prakash Jayaraman, Ali Yavari, and Ahsan Morshed. 2019. "An Industrial IoT Solution for Evaluating Workers' Performance via Activity Recognition." 1393–403.

Gerhart, Barry, and Sara Rynes. 2003. *Compensation: Theory, Evidence, and Strategic Implications*. Sage.

Giermindl, Lisa Marie, Franz Strich, Oliver Christ, Ulrich Leicht-Deobald, and Abdullah Redzepi. 2022. "The Dark Sides of People Analytics: Reviewing the Perils for Organisations and Employees." *European Journal of Information Systems* 31 (3): 410–35.

Gomez-Mejia, Luis R, and Theresa M Welbourne. 1988. "Compensation Strategy: An Overview and Future Steps." *Human Resource Planning* 11 (3): 173–89.

Gong, Yangda, Min Zhao, Qin Wang, and Zhihan Lv. 2022. "Design and Interactive Performance of Human Resource Management System Based on Artificial Intelligence." *PloS One* 17 (1): e0262398.

Gould, Sandy JJ. 2024. "Measuring Work Is Hard. Subcontracting It Won't Help. Explainable AI Won't Help." In *Artificial Intelligence, Labour and Society*. European Trade Union Institute. <https://www.etui.org/publications/artificial-intelligence-labour-and-society>.

Graham, Ronald L. 1966. "Bounds for Certain Multiprocessing Anomalies." *Bell System Technical Journal* 45 (9): 1563–81.

Green, Francis. 2007. *Demanding Work: The Paradox of Job Quality in the Affluent Economy*. Princeton University Press.

Greenberg, Seth Stephens-Davidowitz, Spencer. n.d. "Personality Tests Aren't All the Same. Some Work Better Than Others." *Scientific American*. Accessed January 31, 2025. <https://www.scientificamerican.com/article/personality-tests-arent-all-the-same-some-work-better-than-others/>.

Guest, David E. 2025. "Strengthening Links between HRM Theories, HR Practices and Outcomes: A Proposal to Advance Research on HRM and Outcomes." *Human Resource Management Journal* 35 (1): 319–35.

Hammonds, Keith H. 2005. "Why We Hate HR." *Fast Company* 97 (8): 40–47.

Hiba, Juan Carlos. 1998. *Improving Working Conditions and Productivity in the Garment Industry: An Action Manual*. International Labour Organization.

HRBrain.ai. 2024. *AI in Compensation and Benefits: Predictive Analytics*. January 29. <https://hrbrain.ai/blog/ai-in-compensation-and-benefits-predictive-analytics/>.

ISG. 2023. *2023 Survey on Industry Trends in TR Technology and Service Delivery*. HR Tech Survey. ISG.

Jain, Anant Singh, and Sheik Meeran. 1999. "Deterministic Job-Shop Scheduling: Past, Present and Future." *European Journal of Operational Research* 113 (2): 390–434.

Jaser, Zahira, and Dimitra Petrakaki. 2023. "Are You Prepared to Be Interviewed by an AI?" *Job Interviews*. *Harvard Business Review*, February 7. <https://hbr.org/2023/02/are-you-prepared-to-be-interviewed-by-an-ai>.

Katz, Cindi. 2001. "Vagabond Capitalism and the Necessity of Social Reproduction." *Antipode* 33 (4): 709–28.

Kaufman, Bruce E. 2001. "Human Resources and Industrial Relations: Commonalities and Differences." *Human Resource Management Review* 11 (4): 339–74.

Kaur, Jaspreet, Sheila Arianayagam, and Parmjit Singh. 2020. "The Usage of Social Media Networks as a Tool by Employers to Screen Job Candidates." *The European Journal of Social & Behavioural Sciences* Issue 3 (August). <https://doi.org/10.15405/ejsbs.282>.

Kelley Jr, James E, and Morgan R Walker. 1959. "Critical-Path Planning and Scheduling." 160–73.

Kellogg, Katherine C, Melissa A Valentine, and Angele Christin. 2020. "Algorithms at Work: The New Contested Terrain of Control." *Academy of Management Annals* 14 (1): 366–410.

Kingsley, Sara, Clara Wang, Alex Mikhaleenko, Proteeti Sinha, and Chinmay Kulkarni. 2020. "Auditing Digital Platforms for Discrimination in Economic Opportunity Advertising." arXiv:2008.09656. Preprint, arXiv, August 21. <https://doi.org/10.48550/arXiv.2008.09656>.

Kofman, Ariana Tobin, Ava. 2022. "Facebook Finally Agrees to Eliminate Tool That Enabled Discriminatory Advertising." ProPublica, June 22. <https://www.propublica.org/article/facebook-doj-advertising-discrimination-settlement>.

Kresge, Lisa. 2020. *Data and Algorithms in the Workplace: A Primer on New Technologies*.

Kresge, Lisa. 2023. "Negotiating Workers' Rights at the Frontier of Digital Workplace Technologies in 2023." *UC Berkeley Labor Center*, December 18. <https://laborcenter.berkeley.edu/negotiating-workers-rights-at-the-frontier-of-digital-workplace-technologies-in-2023/>.

Krueger, Alan. 2000. "Economic Scene; The Internet Is Lowering the Cost of Advertising and Searching for Jobs. - The New York Times." *New York Times*, July 20. <https://www.nytimes.com/2000/07/20/business/economic-scene-internet-lowering-cost-advertising-searching-for-jobs.html>.

Krscynski, David, Cody Reeves, Ryan Stice-Lusvardi, Michael Ulrich, and Grant Russell. 2018. "Analytical Abilities and the Performance of HR Professionals." *Human Resource Management* 57 (3): 715–38. <https://doi.org/10.1002/hrm.21854>.

Krzywdzinski, Martin, and Christine Gerber. 2019. "Brave New Digital Work? New Forms of Performance Control in Crowdwork." In *Work and Labor in the Digital Age*, edited by S. P. Vallas and Anne Kovalainen. Emerald (MCB UP).

Kyeremateng-Amoah, Emmanuel, Jackie Nowell, Aurora Luty, Peter SJ Lees, and Ellen K Silbergeld. 2014. "Laceration Injuries and Infections among Workers in the Poultry Processing and Pork Meatpacking Industries." *American Journal of Industrial Medicine* 57 (6): 669–82.

Lambrecht, Anja, and Catherine Tucker. 2019. "Algorithmic Bias? An Empirical Study of Apparent Gender-Based Discrimination in the Display of STEM Career Ads." *Management Science* 65 (7): 2966–81. <https://doi.org/10.1287/mnsc.2018.3093>.

Ledet, Elizabeth, Keith McNutty, Daniel Morales, and Melissa Shandell. 2020. "How to Be Great at People Analytics." *McKinsey and Company: People and Organizational Performance*, October 2.

<https://www.mckinsey.com/capabilities/people-and-organizational-performance/our-insights/how-to-be-great-at-people-analytics>.

Legge, Karen. 1978. *Power, Innovation, and Problem-Solving in Personnel Management*.

Levy, Ferdinand, Gerald L Thompson, and Jerome D West. 1963. *The ABCs of the Critical Path Method*. Harvard University.

Levy, Karen EC. 2015. "The Contexts of Control: Information, Power, and Truck-Driving Work." *The Information Society* 31 (2): 160–74.

Littler, Craig R. 1978. "Understanding Taylorism." *British Journal of Sociology*, 185–202.

Markel, Keith A., Jessica L. Lipson, Alana Mildner Smolow, Keith A. Markel, Jessica L. Lipson, and Alana Mildner Smolow. 2023. "AI and Pay Equity: Positives and Pitfalls for Employers to Consider." Legal Industry. *Reuters*, December 21. <https://www.reuters.com/legal/legalindustry/ai-pay-equity-positives-pitfalls-employers-consider-2023-12-21/>.

Marler, Janet H. 2024. "Artificial Intelligence, Algorithms, and Compensation Strategy: Challenges and Opportunities." *Organizational Dynamics*, 101039.

McCartney, Steven, Caroline Murphy, and Jean Mccarthy. 2021. "21st Century HR: A Competency Model for the Emerging Role of HR Analysts." *Personnel Review* 50 (6): 1495–513.

Mintzberg, Henry. 1987. "The Strategy Concept I: Five Ps for Strategy." *California Management Review* 30 (1): 11–24.

Muldoon, James, Callum Cant, Mark Graham, and Funda Ustek Spilda. 2023. "The Poverty of Ethical AI: Impact Sourcing and AI Supply Chains." *AI & SOCIETY*, 1–15.

OECD. 2021. *The Impact of Artificial Intelligence on the Labour Market: What Do We Know so Far?* OECD Social, Employment and Migration Working Papers No. 256. Vol. 256. OECD Social, Employment and Migration Working Papers. <https://doi.org/10.1787/7c895724-en>.

Paauwe, Jaap, and Corine Boon. 2018. "Strategic HRM: A Critical Review." *Human Resource Management*, 49–73.

Pfeffer, Jeffrey, and Robert I Sutton. 2006. "Evidence-Based Management." *Harvard Business Review* 84 (1): 62.

Raghavan, Manish, Solon Barocas, Jon Kleinberg, and Karen Levy. 2020. "Mitigating Bias in Algorithmic Hiring: Evaluating Claims and Practices." *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, January 27, 469–81. <https://doi.org/10.1145/3351095.3372828>.

Rahman, Hatim A. 2021. "The Invisible Cage: Workers' Reactivity to Opaque Algorithmic Evaluations." *Administrative Science Quarterly* 66 (4): 945–88. <https://doi.org/10.1177/00018392211010118>.

Recalde, María P., and Lise Vesterlund. 2022. "Gender Differences in Negotiation and Policy for Equalizing Outcomes." In *Bargaining: Current Research and Future Directions*, edited by Emin Karagözoğlu and Kyle B. Hyndman. Springer International Publishing. [https://doi.org/10.1007/978-3-030-76666-5\\_21](https://doi.org/10.1007/978-3-030-76666-5_21).

Rousseau, Denise M. 2006. "Is There Such a Thing as 'Evidence-Based Management'?" *Academy of Management Review* 31 (2): 256–69.

Seymour, Diane, and Peter Sandiford. 2005. "Learning Emotion Rules in Service Organizations: Socialization and Training in the UK Public-House Sector." *Work, Employment and Society* 19 (3): 547–64.

Spektor, Franchesca, Sarah E. Fox, Ezra Awumey, et al. 2023. "Charting the Automation of Hospitality: An Interdisciplinary Literature Review Examining the Evolution of Frontline Service Work in the Face of Algorithmic Management." *Proc. ACM Hum.-Comput. Interact.* 7 (CSCW1): 33:1–33:20. <https://doi.org/10.1145/3579466>.

Stabile, Susan J. 2001. "The Use of Personality Tests as a Hiring Tool: Is the Benefit Worth the Cost?" *U. Pa. J. Lab. & Emp. L.* 4: 279.

Suchman, Lucy. 1995. "Making Work Visible." *Communications of the ACM* 38 (9): 56–64.

Tambe, Prasanna, Peter Cappelli, and Valery Yakubovich. 2019. "Artificial Intelligence in Human Resources Management: Challenges and a Path Forward." *California Management Review* 61 (4): 15–42. <https://doi.org/10.1177/0008125619867910>.

Teachout, Zephyr. 2023. "Algorithmic Personalized Wages." *Politics & Society* 51 (3): 436–58.

Terveen, Loren, and David W McDonald. 2005. "Social Matching: A Framework and Research Agenda." *ACM Transactions on Computer-Human Interaction (TOCHI)* 12 (3): 401–34.

Tobin, Julia Angwin, Noam Scheiber, Ariana. 2017. "Dozens of Companies Are Using Facebook to Exclude Older Workers From Job Ads." ProPublica, December 20. <https://www.propublica.org/article/facebook-ads-age-discrimination-targeting>.

Vallas, Steven P, Hannah Johnston, and Yana Mommadova. 2022. "Prime Suspect: Mechanisms of Labor Control at Amazon's Warehouses." *Work and Occupations* 49 (4): 421–56.

Wang, Yingli, Jean-Paul Skeete, and Gilbert Owusu. 2022. "Understanding the Implications of Artificial Intelligence on Field Service Operations: A Case Study of BT." *Production Planning & Control* 33 (16): 1591–607. <https://doi.org/10.1080/09537287.2021.1882694>.

Weaver, Patrick. 2006. "A Brief History of Scheduling." *Melbourne, Australia: Mosaic Project Services Pty Ltd*.

Whang, Steven Euijong, Yuji Roh, Hwanjun Song, and Jae-Gil Lee. 2023. "Data Collection and Quality Challenges in Deep Learning: A Data-Centric AI Perspective." *The VLDB Journal* 32 (4): 791–813. <https://doi.org/10.1007/s00778-022-00775-9>.

Wood, Geoffrey, and Zsuzsanna Kispál-Vitai. 2017. "Strategic Human Resource Management: Concepts, Practices and Trends." *Strategic Human Resource Management: An International Perspective*, 78–105.

Wood, Geoffrey, and Zsuzsanna Kispál-Vitai. 2021. "Concepts, Practices and Trends in SHRM." *Strategic Human Resource Management: An International Perspective*, 74.

Wright, Christopher. 2008. "Reinventing Human Resource Management: Business Partners, Internal Consultants and the Limits to Professionalization." *Human Relations* 61 (8): 1063–86.

Yanadori, Yoshio, and Janet H Marler. 2006. "Compensation Strategy: Does Business Strategy Influence Compensation in High-technology Firms?" *Strategic Management Journal* 27 (6): 559–70.

Zarei, Hamid Reza, Yakov Bart, and Ozlem Ergun. 2023. "Impact of Using a Centralized Matching Process on Nursing Home Staffing." *Geriatric Nursing* 49: 89–93.

Zuboff, Shoshana. 2019. *The Age of Surveillance Capitalism: The Fight for the Future at the New Frontier of Power*. Profile Books.

## Acknowledgements

---

The authors would like to thank Anil Verma, Nikolai Rogovsky and Michael 'Six' Silberman for helpful comments and suggestions.

## ► Advancing social justice, promoting decent work

The International Labour Organization is the United Nations agency for the world of work. We bring together governments, employers and workers to improve the working lives of all people, driving a human-centred approach to the future of work through employment creation, rights at work, social protection and social dialogue.

### Contact details

#### Research Department (RESEARCH)

International Labour Organization  
Route des Morillons 4  
1211 Geneva 22  
Switzerland  
T +41 22 799 6530  
[research@ilo.org](mailto:research@ilo.org)  
[www.ilo.org/research](http://www.ilo.org/research)



I S B N 9789220428627



9 789220 428627