NET WORK LISER

Discussion Paper Series

IZA DP No. 18782

July 2026

Let Me Check on You: Job Quality Under AI and Human Oversight

Milena Nikolova
University of Groningen,
The Brookings Institution, GLO,
IZA@LISER and Netspar

IZA@LISER NETWORK
ISSN: 2365-9793

Let Me Check on You: Job Quality Under AI and Human Oversight*

Abstract

This paper provides the first causal evidence on how Artificial Intelligence (AI)-based workplace safety systems shape perceived job quality. I conducted a pre-registered vignette experiment with a nationally representative sample of 2,172 Dutch adults who evaluated otherwise identical workplaces introducing one of three safety systems: i) human supervisors, ii) AI-only monitoring, or iii) hybrid AI-human supervision. Compared with human supervision, both AI-only and hybrid systems reduced perceived job satisfaction, work meaningfulness, and perceived social value of the job. Contrary to expectations, combining AI with human supervisors did not mitigate these negative effects. Respondents also viewed AI-based systems as less respectful of workers' privacy and dignity, despite viewing them as effective as human supervisors. Perceived fair wages changed little across conditions. These findings suggest that AI can influence work not only by improving safety but also by reducing important non-pecuniary dimensions of job quality, highlighting that the welfare consequences of workplace AI extend beyond productivity and accident prevention.

JEL classification

I39, J01, J28, O33

Keywords

Artificial Intelligence (AI), safety systems, survey experiment, work meaningfulness, job quality

Corresponding author

Milena Nikolova
m.v.nikolova@rug.nl

* *Acknowledgments*: This project received funding from the Lloyd's Register Foundation under the Turning the 2023 World Risk Poll into Action scheme (grant agreement no. W100006). The survey experiment was developed in collaboration with Viliana Milanova (University of Groningen). I thank Lieke Vooringholt, Maite Laméris, Daniele Marchesi, and Marco Angrisani for valuable feedback on the experimental design. Any remaining errors are my own.

Declaration of AI Use: The author used OpenAI's ChatGPT (versions 5.4 and 5.5) for language editing, copyediting, and improving the clarity and readability of the manuscript. The author also used OpenAI Codex and ChatGPT (version 5.4) to assist with coding for data analysis. All AI-generated text and code were reviewed, verified, and revised by the author. The author takes full responsibility for the manuscript's content, analyses, code accuracy, and the conclusions reported in the paper.

1. Introduction

Artificial Intelligence (AI) is an ever-evolving technological frontier (Berente et al., 2021), with the capability to process large volumes of data, recognize patterns, and generate real-time insights. In the occupational health and safety domain, AI comprises data-driven systems that perform pattern recognition tasks, enabling prediction, assessment, and recommendations to reduce accidents and improve safety.

AI can be especially effective in detecting workplace hazards, monitoring the correct use of protective gear and equipment, and predicting accident rates (La Torre et al., 2026). Machine-learning-based systems are particularly good at analyzing historical accident data and discovering trends that can help prevent further accidents. AI programs can identify and assess exposure to hazardous gases and substances, and detect equipment malfunctions that human observers may miss (El-Helaly, 2024). Similarly, deep learning techniques help optimize the design of personal protective equipment (PPE) (La Torre et al., 2026). For example, wearable devices equipped with AI, such as smart helmets, monitor vital signs, steps taken, fatigue levels, and can detect harmful gases or dangerous working conditions (El-Helaly, 2024). In many hazardous occupations, including agriculture, construction, and mining, smart PPE augments traditional protective equipment with embedded electronics and AI capabilities that support real-time monitoring and risk prevention (El-Helaly, 2024). AI systems can also use voice recognition and natural language processing to detect risks related to workplace violence and harassment (El-Helaly, 2024).

Yet, in collecting and processing this information, the use of AI systems also raises concerns about privacy and security, surveillance, and workers' loss of autonomy (Probst et al., 2026). AI systems have algorithmic biases due to how they are trained, which can raise concerns about discrimination and unfair treatment in the workplace. In many settings, workplace interactions are complex, raising questions about the sufficiency of AI capabilities to handle these complexities, the likelihood of errors, and how much AI systems should be trusted (La Torre et al., 2026). Like other automation technologies, AI safety systems also may induce fear of job displacement (Küsters & Schneider, 2025; Mokyr et al., 2015) and discourage safety professionals from continuing to invest in their skills (El-Helaly, 2024). In most workplace settings, AI safety systems do not operate independently but rather serve a supportive role by examining accident reports, generating risk assessments, and monitoring the proper use of protective equipment (La Torre et al., 2026). The prospect of full decision automation and deployment of AI systems is increasing, as these systems become more sophisticated.

Against this backdrop, evidence on the impact of AI on workplace health and safety and worker well-being remains very limited (De Simone, 2026; Jetha et al., 2025; Toole et al., 2025). Yet, worker well-being is important for employers and the economy, because it correlates with productivity, effort, health, and turnover (Cassar & Meier, 2018; Clark et al., 2012; Clark, 2001; Cnossen & Nikolova, 2025; De Neve et al., 2013; Nikolova & de Wit, 2026).

Against this backdrop, the majority of scientific work on technological change and worker health and safety has focused on industrial robots. Specifically, robots reduce workplace injuries through reducing physical strain (De Simone et al., 2025; Gihleb et al., 2022; Gunadi & Ryu, 2021). Industrial robots also increase psychosocial risks through increased job insecurity (Schwabe & Castellacci, 2020), increased monotonous work (Nikolova et al., 2025), and worsened mental health (Abeliansky et al., 2024). Industrial robots have also reduced the meaningfulness and autonomy of European workers (Nikolova et al., 2024).

The limited available evidence from nationally representative studies in Finland and Germany suggests that AI does not yet have large impacts on workplace well-being (Bryson et al., 2026; Giuntella et al., 2025). Giuntella et al. (2025) also show that AI may lead to a lower physical burden and small improvements in health satisfaction in Germany (Giuntella et al., 2025). Finally, using survey data from nearly 5,000 UK employees, Soffia et al. (2024) find that exposure to AI software is associated with greater perceived salary progression, career prospects, autonomy, and work meaningfulness, but also with increased job insecurity.

Yet existing studies provide no causal evidence on whether introducing AI into health and safety systems changes perceptions of job satisfaction and work meaningfulness. To shed light on this question, I present the first causal evidence from a vignette experiment. Using data from a nationally representative sample of 2,172 Dutch adults, I find that respondents perceive AI-based workplace safety systems as reducing several dimensions of job quality relative to traditional human supervision. Compared with human safety supervisors, fully AI-based systems lower perceived job satisfaction by 0.15 standard deviations, work meaningfulness by 0.11 standard deviations, and the perceived social value of the job by 0.10 standard deviations. Hybrid systems that combine AI and human supervisors produce nearly identical effects, reducing job satisfaction by 0.12 standard deviations, meaningfulness by 0.11 standard deviations, and social value by 0.10 standard deviations. By contrast, the effects on what the perceived fair wages that workers should receive for working on this job are small and inconsistent. Beyond job quality, respondents also perceive AI-only and hybrid safety systems as equally effective in improving health and safety but less respectful of workers' privacy and dignity than traditional human supervision. Overall, the findings suggest that while AI has a promising role in improving health and safety in the future, workers perceive AI-enabled monitoring systems as diminishing important non-pecuniary aspects of work, including satisfaction, meaning, and the job's societal value.

The Netherlands provides an informative setting for studying the work well-being implications of AI-based workplace safety systems because it combines relatively high levels of workplace safety[1] with comparatively high rates of AI adoption. Recent evidence indicates that more than one-third of Dutch workers use generative AI at work, placing the Netherlands among the leading adopters in Europe (Bick et al., 2026). At the same time, data from the 2023 World Risk Poll show that workplace risks remain relevant in the Netherlands. Figure 1 shows that 27% of Dutch workers report being worried that their work could cause serious harm, while 8% report

[1] The Netherlands had one of the lowest rates of fatal workplace accidents in the EU in 2023, with fewer than one fatal accident per 100,000 employees. By contrast, its incidence of non-fatal workplace accidents ranked around the middle of the EU distribution. In 2023, the Netherlands reported nearly 85,000 non-fatal workplace accidents, corresponding to roughly 1,000 accidents per 100,000 workers.

having personally experienced serious workplace harm during the previous two years. Although these figures are substantially lower than the corresponding global averages (LRF, 2024), they suggest that workplace safety remains an important concern even in a relatively safe and technologically advanced labor market.

Figure 1: Workplace health and safety in the Netherlands

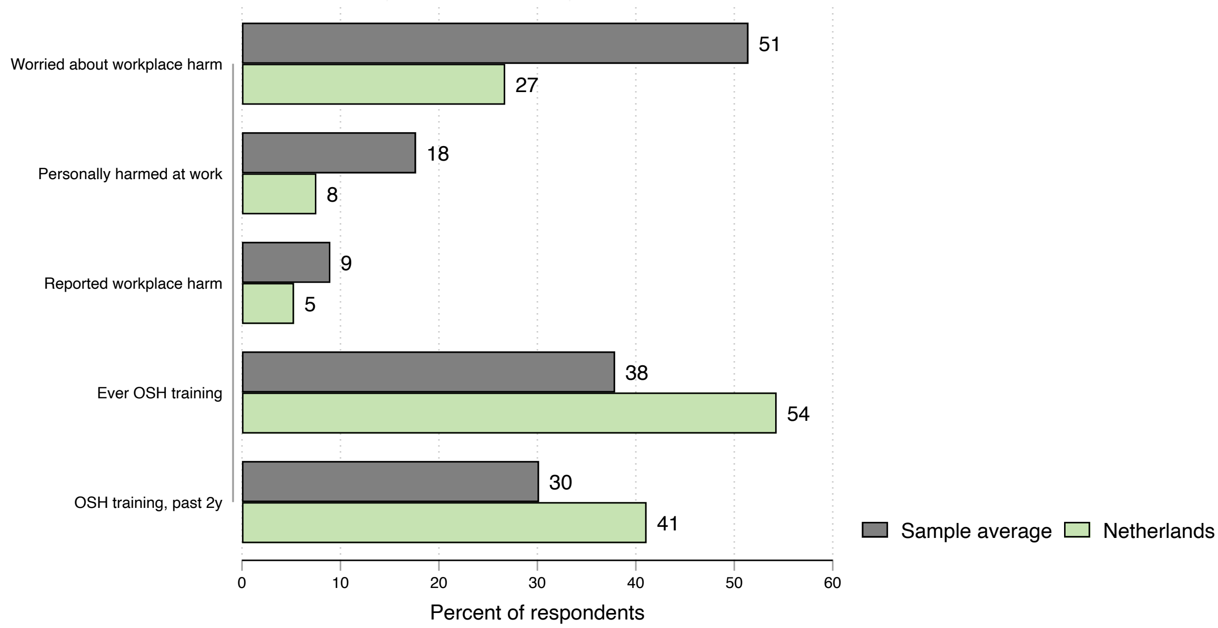

Notes: Author's calculations based on the World Risk Poll (2024). The figure reports weighted shares for respondents in the labor force, defined as individuals who are employed or unemployed.

This paper contributes to the emerging literature on AI and workplace well-being by providing the first causal evidence on how AI-based workplace safety systems affect perceptions of job quality. The findings have important implications for the future design of workplace technologies. Discussions of AI in occupational health and safety often focus on efficiency gains and accident prevention (Toole et al., 2025). However, jobs provide not only income but also meaning, satisfaction, and a sense of social contribution (Cassar & Meier, 2018; Cnossen & Nikolova, 2025; Hendriks & Cnossen, 2026; Nikolova & Cnossen, 2020; Rosso et al., 2010). Understanding whether AI-based monitoring systems alter these non-pecuniary aspects of work is therefore essential for evaluating the broader welfare consequences of workplace automation.

2. Conceptual Framework and Hypotheses

This paper studies how the method of being supervised and monitored for health and safety reasons influences perceptions of job quality in terms of work meaningfulness, job satisfaction, the societal purpose of the job, and the appropriate or "fair" wages that workers deserve.

2.1. Defining Work Meaningfulness, Societal Purpose, Job Satisfaction, and Fair Wages

Meaningful work refers to the extent to which individuals perceive their work as significant, worthwhile, and contributing to a broader personal or societal purpose (Cnossen & Nikolova, 2025; Hendriks & Cnossen, 2026; Nikolova & Cnossen, 2020; Rosso et al., 2010). Work meaningfulness depends on the extent to which individuals can satisfy their innate psychological needs from working, namely competence (i.e., the feeling of mastery and applying one's own skills), autonomy (i.e., having agency and a sense of control over one's work), and relatedness (i.e., the relationships with clients, co-workers, and superiors). Together, these factors explain 60% of the variation in work meaningfulness, with relatedness being the most important one (Nikolova & Cnossen, 2020). Work meaningfulness can be an important determinant of labor supply (Ariely et al., 2008; Cassar & Meier, 2018; Non et al., 2022), effort (Cnossen & Nikolova, 2025), and later retirement and training participation (Nikolova & Cnossen, 2020).[2]

Some studies distinguish between work meaningfulness arising from societal and non-social impact (Burbano et al., 2024). Societal impact refers to the job's contributions to society and community more generally, such as working in healthcare, protecting the environment, or working with underprivileged groups, among others. The prosocial mission of an organization, such as a public organization or a nonprofit, also matters to workers (Burbano et al., 2024; Burbano, 2016; Cassar & Meier, 2018; Non et al., 2022). In fact, several papers on work meaningfulness define and measure it solely in terms of its pro-social impact and societal usefulness (Dur & Van Lent, 2019; Kesternich et al., 2021; Schouwer et al., 2025). The non-social impact related to work meaningfulness, furthermore, concerns personal significance, pride, and the degree to which work satisfies the innate needs for autonomy, competence, and relatedness (Burbano et al., 2024).

In principle, a job need not have societal significance to be meaningful (Burbano et al., 2024; Hendriks & Cnossen, 2026; Rosso et al., 2010). This is why in this paper, I treat the societal importance of the job as a separate but related construct to work meaningfulness. It is ultimately an empirical question of whether using both constructs separately produces similar or different results in my setting.

Job satisfaction is a related but distinct aspect of perceived job quality at work (Nikolova & Cnossen, 2020). Work meaningfulness captures whether work feels significant and purposeful. Job satisfaction reflects a broader evaluation of one's job, including pay, working conditions, expectations, and rewards. Consequently, an employee can report high job satisfaction while finding the work relatively meaningless, or experience highly meaningful work despite dissatisfaction with pay or working conditions. Job satisfaction predicts workplace behaviors, such as job quits (Böckerman & Ilmakunnas, 2009; Clark et al., 2012; Clark, 2001; Cornelißen, 2009; Lévy-Garboua et al., 2007; Nikolova & de Wit, 2026).

[2] Hendriks and Cnossen (2026) provide a nuanced summary of the causes and consequences of meaningful work.

Finally, fair wages refer to the level of compensation that workers deem appropriately compensates their level of effort (Akerlof & Yellen, 1990) and given the working conditions they face (Schouwer et al., 2025). This concept is also closely related to the theory of compensating wage differentials (Rosen, 1986), which posits that workers receive higher compensation for adverse working conditions, such as noise, hazards, heat, poor ergonomics, etc. In this sense, being supervised by humans or AI could bring advantages and disadvantages.

2.2. Hypotheses Development

Many occupations and workplaces have employees who coordinate, supervise, monitor, and support work processes. Human supervisors do not merely enforce rules and monitor performance. They also create opportunities for social interaction, guidance, and interpersonal relationships, thereby contributing to workers' sense of relatedness. Prior research shows that managing others can enhance work meaningfulness (Nikolova et al., 2023), and that may be due to the social aspects of supervision. Having good relationships with colleagues, supervisors, and subordinates creates a sense of belonging in the organization and facilitates camaraderie and social connections among workers. Because relatedness is a key determinant of meaningful work (Nikolova & Cnossen, 2020), workers may perceive jobs as less meaningful when functions traditionally performed by humans are delegated to technological systems. Human supervision can also signal interpersonal responsibility and accountability, reinforcing the perception that work contributes to a broader social process. Evidence from previous research on robots suggests that, because technologies standardize the production process, there is less product-quality heterogeneity and less need for supervision (Dixon et al., 2021). Similarly, robotization leads to a decrease in relatedness at work (Nikolova et al., 2024) and in the social tasks workers perform at work (Nikolova et al., 2025). Furthermore, an AI-based safety system may be invoking issues related to algorithmic management, such as loss of autonomy due to standardized tasks, time pressure, and the feeling of eeriness about being watched, and a feeling of insecurity related to job loss (Noponen et al., 2024; Riso et al., 2025).

The introduction of AI systems and the implicit replacement of some human oversight by these systems may affect meaningfulness, job satisfaction, and the perceived social importance of the job if individuals view AI oversight as reducing the uniquely human component of organizational activity. This argument aligns with recent evidence showing that AI substitution can reduce perceptions of meaningful work even when workers themselves are not directly replaced, but some parts of their work process are (Nikolova et al., 2026; Sadeghian & Hassenzahl, 2022; Zhang & Amos, 2026). The effect emerges because AI changes how individuals interpret the significance of human effort within the production process. As AI assumes responsibilities previously associated with human expertise and judgment, workers may perceive their own role as less important, less socially valuable, and less intrinsically rewarding. In this scenario, if AI systems diminish perceived job quality, individuals may believe that workers should be compensated with higher wages. In that sense, having AI monitoring for safety purposes may be a disamenity.

In principle, the adoption of AI-based safety systems may increase perceived job quality when workers view AI as a partner in maintaining their health and safety rather than as a tool for monitoring or control. In this role, AI augments human capabilities by identifying workplace hazards, detecting signs of fatigue, and recommending appropriate actions. Workers retain responsibility for interpreting and implementing these recommendations, which preserves their sense of agency and control. Consequently, they may perceive the AI system as a mentor that supports safer and more effective work (Noponen et al., 2024). This perception may also strengthen the meaningfulness of work and job satisfaction by enhancing workers' sense of competence. By helping workers apply their skills effectively while reducing safety risks, AI enables them to perform their jobs to the best of their abilities and with greater confidence. In this setting, the perceived fair wages may not change and may even decrease if AI is perceived to be a net benefit. Furthermore, the perceived societal importance of the job is unlikely to change if AI serves an augmenting function, unless participants equate the societal importance of a job with work meaningfulness.

The effects of AI-based safety systems are unlikely to depend solely on whether AI is present in the workplace, but also on the extent to which it replaces or complements human supervision. A hybrid system, in which AI supports a human safety supervisor, represents an intermediate form of oversight. In this setting, AI provides information, detects hazards, and generates recommendations, while the human supervisor retains responsibility for interpreting this information, making final decisions, and interacting with workers. As a result, workers continue to experience interpersonal guidance and accountability while benefiting from AI's analytical capabilities.

From the perspective of meaningful work, the hybrid system may preserve the sense of relatedness associated with human supervision while strengthening competence by providing more accurate information and decision support. Likewise, concerns related to algorithmic management, such as loss of autonomy, surveillance, or dehumanization, may become less salient because workers remain accountable to a human supervisor rather than to an algorithm. Consequently, the effects of the hybrid system should lie between those of purely human and purely AI-based oversight. Relative to human supervision, the hybrid system may reduce meaningfulness and job satisfaction only modestly, if at all. Relative to AI-only supervision, however, the hybrid system should result in higher perceived meaningfulness, job satisfaction, and social value, while reducing any compensating wage premium associated with AI oversight.

In the empirical part of the paper, I examine whether respondents perceive AI-based safety systems as reducing or enhancing job quality and whether hybrid AI-human supervision mitigates any adverse perceptions associated with AI-only supervision.

3. Experimental Design

I use a between-subjects vignette experiment to examine how different workplace safety systems affect perceptions of worker safety, health risks, and job quality. Vignette experiments present individuals with realistic scenarios that they evaluate to elicit their preferences, attitudes, and potential behaviors. I employ a vignette experiment that presents respondents with a realistic workplace scenario while varying a key attribute of interest, namely the type of workplace safety system. This design allows us to isolate the causal effect of AI-based, human-based, and hybrid safety monitoring systems while maintaining a manageable cognitive burden for respondents (Stantcheva, 2023). Previous validation research finds that responses in vignette experiments often align closely with real-world decision-making (Hainmueller et al., 2015), although alternative methods such as paired designs tend to exhibit greater external validity.

At the beginning of the survey, all respondents read a baseline vignette describing a 40-year-old assembly-line worker employed in a large manufacturing plant. The vignette stated that the company complies with occupational safety regulations, provides protective equipment and safety training, and maintains a positive and collegial work environment. Before moving on to the outcome questions, respondents completed a comprehension check asking them to identify the safety measures described in the vignette. The vignette text was displayed on the same screen as the comprehension check question, allowing respondents to reread the vignette.

Participants who answered incorrectly were encouraged to reread the scenario before continuing and were given another chance to answer the question. In line with LISS panel procedures, we allowed all respondents to complete the questionnaire. Respondents then evaluated the job on several dimensions, including safety, risk, stress, physical demands, the likelihood of workplace accidents and physical and mental health problems, job satisfaction, meaningfulness, social usefulness, and appropriate pay. During the survey, respondents always had the option to review the vignette text.

These baseline assessments capture perceptions of the job before information about the introduction of the new safety system was provided. The outcomes I focus on in this paper include i) job satisfaction, ii) work meaningfulness, iii) perceptions of whether the work is socially important, and the fair wage that such a worker should receive.

Next, respondents read a description of a newly introduced workplace safety system that corresponded to their randomly assigned treatment condition. In all three conditions, the system continuously monitors workplace risks, detects hazards such as malfunctioning machinery, missing protective equipment, and signs of worker fatigue, and provides immediate safety recommendations. The vignettes had similar wording across conditions. Respondents were told that all three systems perform the same safety functions and pursue the same objective of preventing workplace accidents and protecting worker health. The vignettes varied only in one feature: whether AI, trained human safety supervisors, or a combination of both carry out the monitoring and provide recommendations (e.g., taking a break). The experimental manipulation, therefore, varies the source of workplace safety monitoring while holding constant the system's stated purpose: the real-time detection of workplace risks and hazards. Random assignment allows for the interpretation of differences in respondents' evaluations as the causal effect of the type of safety system. After reading the treatment vignette, respondents completed a second comprehension check that assessed their understanding of the safety system. They had two chances to answer the comprehension check correctly and could continue answering the survey even if they did not provide the correct answers.

After the treatment, respondents evaluated the job again after the implementation of the new safety system using the same measures as in the baseline assessment. They again had access to the vignette text if they needed it. Finally, respondents evaluated the safety system by reporting how much they trusted it to prevent accidents and improve worker health, how effective they considered it, and how well it respected workers' privacy and dignity.

The main advantage of this paper's research design is that it provides causal evidence on how individuals perceive alternative workplace safety technologies while preserving a realistic occupational setting. The main limitation is that respondents evaluate hypothetical workplace scenarios rather than interact with actual safety technologies in real workplaces. Consequently, the results capture perceptions and expectations regarding AI-enabled safety systems rather than their actual effects on workplace safety and worker well-being.

4. Procedures, Data, and Variables

The experiment was preregistered in the American Economic Association Randomized Controlled Trial Registry before data collection began (AEARCTR-0016911). The study also received ethical approval from the Institutional Review Board of the Faculty of Economics and Business at the University of Groningen (protocol number FEB-20251001-01598).

Data collection took place in October and November 2025 through the Dutch Longitudinal Internet Studies for the Social Sciences (LISS) panel, a nationally representative probability-based panel administered by CenterData at the University of Tilburg (Das & Knoef, 2019; Scherpenzeel, 2011). The LISS sample is a true probability sample of households from the Dutch population registry, comprising 5,000 households and 7,500 individuals. Panel members answer questions monthly and receive compensation for doing so.

Fully 2,775 panel members aged 16 and older received an invitation to complete our survey, and 2,187 did so, yielding a response rate of 78.8%. We wrote the questionnaire first in English and then translated it into Dutch. The LISS team checked the questionnaire before programming it. We excluded 15 respondents who did not complete the questionnaire, resulting in a final sample of 2,172 respondents. Of these, 733 respondents received the AI treatment, 723 received the human-only treatment, and 716 received the hybrid treatment.

The four outcome variables are (i) job satisfaction, (ii) work meaningfulness, (iii) perceived social value, and (iv) perceived fair wages (see Appendices A and B). Before participants read the vignette, they rated the first three outcomes by answering the question: "Based on the description, to what extent do you think this worker...?" The three sub-questions asked whether the worker "is satisfied with the job," "finds the job meaningful," and "finds the job socially useful." Participants answered each question on a seven-point scale ranging from 1 ("Not at all") to 7 ("Very much").

Respondents hey also reported the worker's perceived fair monthly gross wage by answering the following question: "Based on the description, what do you think is an appropriate monthly gross salary for this worker? Assume a 40-hour work week and a salary excluding bonuses." Participants selected one of seven response categories in €500 increments, ranging from less than €1,500 to €4,000 or more.

After respondents read the treatment vignette, they answered the same questions using the wording: "After the implementation of the new safety system, to what extent do you think this worker...?" Respondents then indicated the extent to which they thought the worker "is satisfied with the job," "finds the job meaningful," and "finds the job socially useful" on the same seven-point scale. They also reported the worker's perceived fair monthly gross wage using the same salary question and response categories as before the treatment.

In the regression analyses, I standardize all outcome variables to have a mean of zero and a standard deviation of one to facilitate comparison across outcomes. For the fair wage measure, I first convert the categorical response into a continuous variable using the midpoint of each wage interval and then standardize the resulting measure.[3] I also construct change variables defined as the difference between the unstandardized post-treatment and pre-treatment values of each outcome. The change variable itself is standardized.

Furthermore, I use information on respondents' socio-demographic characteristics drawn from the LISS panel. I control for age; a female indicator; a higher-education indicator equal to one for respondents with higher vocational (HBO) or university (WO) education, with a separate category for missing income; a married indicator equal to one for married respondents and zero otherwise; a working indicator equal to one for respondents in paid work, family work, or self-employment; an urban-residence indicator, with a separate category for missing income; an indicator for living alone; an indicator for children in the household; and indicators for terciles of personal net income, with a separate category for missing income.

I also linked the survey data to the LISS Personality modules to control for respondents' Big Five traits. Our primary source was the 2025 LISS Personality file (cp25q), fielded in May-June 2025. When a respondent was missing a 2025 personality item, we filled that item using the corresponding response from the 2024 Personality file (cp25q), fielded in May-June 2024. Using the IPIP Big-Five battery (items 020-069), I reverse-coded negatively worded items and then constructed indices for extraversion, agreeableness, conscientiousness, emotional stability, and imagination/intellect, following the Goldberg Big 5 personality scales (Goldberg, 1992) by summing the items. All personality variables are standardized to have a mean of 0 and a standard deviation of 1.

[3] Two individuals noted in the comments to the survey that they put the wrong value for the pre-treatment fair wages and indicated the correct value that they wanted to select. Therefore, I adjusted their responses.

Table 1 shows that random assignment produced comparable treatment groups. The three groups exhibit similar demographic characteristics. Across the balance tests, only one income category differs significantly across treatment conditions. Most importantly, the treatment groups do not differ significantly on any of the baseline outcome measures. Although randomization achieved good overall balance, the main regressions nevertheless control for baseline outcome levels to improve statistical precision. In additional specifications, we also control for demographic characteristics and income terciles.

Table 1: Balancing of covariates, by treatment status

	Human (N=723)	AI (N=733)	Hybrid (N=716)	Difference *p*-value
	Panel A: Respondent characteristics			
Age	55.041	56.658	54.869	0.110
	(18.817)	(17.725)	(18.043)	
Female	0.508	0.510	0.532	0.593
	(0.500)	(0.500)	(0.499)	
High school or less	0.560	0.566	0.566	0.968
	(0.497)	(0.496)	(0.496)	
Higher education	0.434	0.430	0.426	0.950
	(0.496)	(0.495)	(0.495)	
Education missing	0.006	0.004	0.008	0.586
	(0.074)	(0.064)	(0.091)	
Married	0.535	0.512	0.499	0.369
	(0.499)	(0.500)	(0.500)	
Working	0.483	0.484	0.494	0.891
	(0.500)	(0.500)	(0.500)	
Rural	0.306	0.295	0.286	0.723
	(0.461)	(0.456)	(0.452)	
Urban	0.693	0.700	0.711	0.754
	(0.462)	(0.459)	(0.454)	
Urban missing	0.001	0.005	0.003	0.399
	(0.037)	(0.074)	(0.053)	
Single-person household	0.243	0.228	0.249	0.624
	(0.429)	(0.420)	(0.433)	
Children in household	0.315	0.340	0.352	0.323
	(0.465)	(0.474)	(0.478)	
Income tertile 1	0.318	0.286	0.344	0.063
	(0.466)	(0.452)	(0.475)	
Income tertile 2	0.295	0.362	0.311	0.019
	(0.456)	(0.481)	(0.463)	
Income tertile 3	0.325	0.299	0.297	0.446
	(0.469)	(0.458)	(0.457)	
Income missing	0.062	0.053	0.047	0.469
	(0.242)	(0.225)	(0.213)	

Panel B: Baseline outcomes				
Baseline job satisfaction	4.867	4.802	4.830	0.526
	(1.079)	(1.122)	(1.101)	
Baseline meaning	4.552	4.584	4.573	0.864
	(1.128)	(1.174)	(1.135)	
Baseline social value	4.350	4.367	4.402	0.702
	(1.184)	(1.238)	(1.213)	
Baseline fair wage (EUR)	2,812.241	2,869.372	2,826.117	0.135
	(564.042)	(571.171)	(556.087)	

Notes: Baseline means are reported by treatment group, with standard deviations in parentheses. The final column reports p-values from tests of equality of means across the AI, Human, and Hybrid conditions.

5. Econometric Strategy

I estimate the effect of alternative workplace safety systems on perceived job quality using Ordinary Least Squares (OLS) regressions of post-treatment outcomes on treatment indicators and the corresponding baseline outcome measure. The main specification is:

$$Y_{i1} = \alpha + \beta_1 AI_i + \beta_2 Hybrid_i + \gamma Y_{i0} + \epsilon_i,$$

where Y_{i1} denotes the post-treatment outcome for individual *I*, Y_{i0} is the corresponding pre-treatment outcome, and AI_i and $Hybrid_i$ are indicators for assignment to the AI and hybrid treatment conditions, respectively. The omitted category is the human-supervisor condition. Accordingly, β_1 and β_2 capture the effects of the AI and hybrid conditions relative to the human condition.

I estimate this specification separately for four outcomes: meaning, job satisfaction, social value, and fair wages. Because treatment is randomly assigned, differences across treatment arms have a causal interpretation. Including the baseline value of the outcome improves precision and helps account for potential imbalances across groups.

I next estimate an extended specification that adds a vector of respondent characteristics:

$$Y_{i1} = \alpha + \beta_1 AI_i + \beta_2 Hybrid_i + \gamma Y_{i0} + \delta' X_i + \epsilon_i,$$

where X_i includes age, a female indicator, higher education, marital status, employment status, urban residence, living alone, and whether children are present in the household. In addition, we include income tercile indicators based on personal net income, as well as a separate category for respondents with missing income data. I estimate all regressions using heteroskedasticity-robust standard errors.

I examine treatment heterogeneity by age, gender, employment status, education, and income, based on interactions between the AI and Hybrid treatment variables and their respective indicators for these socio-demographic characteristics. I also examine heterogeneity based on personality traits.

To assess the robustness of the main conclusions, we also estimate change-score specifications of the form:

$$\Delta Y_i = \alpha + \beta_1 AI_i + \beta_2 Hybrid_i + \delta' X_i + \epsilon_i,$$

where $\Delta Y_i = Y_{i1} - Y_{i0}$.

Finally, I assess the robustness of the main findings in several ways. First, I re-estimate the main models after restricting the sample to respondents who passed both comprehension checks and had survey completion times exceeding the 5^{th} percentile of the response-time distribution. This restriction excludes 218 respondents, leaving an analysis sample of 1,954. Second, I augment the main specification with controls for respondents' evaluations of the questionnaire. Finally, I control for the Big Five personality traits to account for stable individual differences in respondents' evaluations of hypothetical jobs. Across all robustness checks, the estimated treatment effects remain highly stable.[4]

6. Results

6.1. Main Results

Table 2 presents the effects of alternative workplace safety systems on perceptions of job quality. Across all specifications, respondents evaluate jobs that use AI-based safety systems less favorably than jobs that rely exclusively on human safety supervisors. Relative to the human-supervisor condition, the AI-based safety system reduces perceived job satisfaction by 0.15 standard deviations (Column (1)), work meaningfulness by 0.11 standard deviations (Column (2)), and perceived social value by 0.10 standard deviations (Column (3)). As expected, including the demographic controls has little effect on the estimated effects, indicating that observable respondent characteristics do not drive the results.

The effect sizes are rather small but are in line with the limited literature on the impact of AI on well-being more generally. For example, using a sample of German workers, Guintella et al. (2025) find that working with AI reduces job satisfaction by approximately 0.05 standard deviations, an effect that is roughly three times smaller than the estimates reported here. Direct comparisons between the two studies are impossible, however. The two studies differ in several important respects, including the country context (Germany versus the Netherlands), the time period (2000-2020 versus 2025), the measure of AI exposure (working with AI versus a hypothetical exposure to a specific workplace AI safety system), and the empirical methodology

[4] I also restrict the analysis sample to respondents of working age by dropping those older than 66. The results, available upon request, do not change substantively.

(difference-in-differences versus a vignette experiment).

Table 2 also shows that the hybrid safety system generates effects that closely resemble those of the fully AI-based system. Relative to the human-supervisor condition, the hybrid system reduces perceived job satisfaction by 0.13 standard deviations, work meaningfulness by 0.11 standard deviations, and perceived social value by 0.10 standard deviations. These estimates differ only marginally from those of the AI-only system. Thus, the presence of a human supervisor does not appear to mitigate the negative perceptions associated with the introduction of AI into workplace safety. The results support the prediction that AI-based safety systems reduce perceived job quality relative to human supervision. However, they do not support the expectation that hybrid supervision attenuates these negative effects

Furthermore, I initially expected the treatment to have little effect on perceived social value because the safety system does not change the worker's core tasks or the job's perceived contribution to society. Contrary to this expectation, respondents assign lower social value to jobs that use AI-based or hybrid safety systems. One possible explanation is that respondents do not distinguish sharply between social value and work meaningfulness when evaluating hypothetical jobs. This interpretation receives some support from the data, as the post-treatment correlation between the two measures equals 0.82. Rather than evaluating the job's objective contribution to society, respondents may incorporate their broader assessment of the job's meaningfulness into both measures.

Finally, the evidence regarding the impact of different systems on respondents' perceived level of compensation (i.e., fair pay) is considerably weaker. The AI treatment has no statistically significant effect on perceived fair wages. The hybrid treatment produces a small, marginally statistically significant negative effect. Overall, the results provide little support for a compensating wage differential mechanism whereby workers would require higher wages to offset perceived disadvantages associated with AI-based workplace technologies.

Table 2: Post-treatment outcomes

	(1) Job satisfaction	(2) Meaning	(3) Social value	(4) Fair wage	(5) Job satisfaction	(6) Meaning	(7) Social value	(8) Fair wage
AI vs. Human	-0.154^{***}	-0.110^{***}	-0.095^{***}	-0.016	-0.163^{***}	-0.117^{***}	-0.105^{***}	-0.016
	(0.045)	(0.041)	(0.036)	(0.022)	(0.045)	(0.041)	(0.037)	(0.022)
Hybrid vs. Human	-0.120^{***}	-0.108^{***}	-0.098^{***}	-0.042^{*}	-0.125^{***}	-0.112^{***}	-0.103^{***}	-0.041^{*}
	(0.043)	(0.039)	(0.035)	(0.022)	(0.043)	(0.039)	(0.036)	(0.022)
Controls	No	No	No	No	Yes	Yes	Yes	Yes
Baseline outcome	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	2,172	2,172	2,172	2,172	2,172	2,172	2,172	2,172
R-squared	0.287	0.422	0.531	0.828	0.293	0.431	0.536	0.830

Notes: All dependent variables are standardized to have a mean of 0 and a standard deviation of 1. Controls include age, gender, and indicators for: higher education degree, married, working status, urban residence, living alone, children in the household, and high-income tercile. Robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

6.2. Robustness Checks

First, I examine whether the main findings depend on excluding respondents who failed one or both comprehension checks and those who completed the survey unusually quickly. Specifically, I restrict the sample to respondents who passed both comprehension checks and had completion times exceeding the 5th percentile of the response-time distribution. This restriction reduces the sample from 2,172 to 1,954 respondents (90% of the original sample).

Table 3 shows that the main conclusions remain unchanged. Relative to the human-supervisor condition, both the AI-based and hybrid safety systems reduce perceived job satisfaction, job meaning, and social value. The estimated coefficients are similar in magnitude to those reported in the main specification, indicating that the results do not depend on inattentive respondents or unusually short completion times.

The evidence for fair pay remains weak. The AI treatment has no statistically significant effect on perceived fair wages. The hybrid treatment produces a small negative effect of approximately 0.05 standard deviations. Using this analysis sample's standard deviation of perceived fair wages (€543), this estimate corresponds to roughly €27 per month. Although statistically significant, the magnitude of the effect is economically small.

Table 3: Robustness check: attention-check sample

	(1) Job satisfaction	(2) Meaning	(3) Social value	(4) Fair wage	(5) Job satisfaction	(6) Meaning	(7) Social value	(8) Fair wage
AI vs. Human	-0.171***	-0.115***	-0.097**	-0.010	-0.179***	-0.121***	-0.105***	-0.010
	(0.048)	(0.043)	(0.038)	(0.020)	(0.048)	(0.043)	(0.038)	(0.020)
Hybrid vs. Human	-0.136***	-0.102**	-0.082**	-0.045**	-0.143***	-0.108***	-0.087**	-0.046**
	(0.045)	(0.041)	(0.037)	(0.021)	(0.046)	(0.041)	(0.037)	(0.021)
Controls	No	No	No	No	Yes	Yes	Yes	Yes
Baseline outcome control	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1,954	1,954	1,954	1,954	1,954	1,954	1,954	1,954
R-squared	0.261	0.421	0.546	0.850	0.269	0.431	0.552	0.852

Notes: All dependent variables are standardized to have a mean of 0 and a standard deviation of 1. Sample restricted to respondents who passed both attention checks and whose completion time was at or above the 5th percentile. Controls include age, gender, and indicators for: higher education degree, married, working status, urban residence, living alone, children in the household, and high-income tercile. Robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Third, Table 4 re-estimates the treatment effects using change scores, defined as post-treatment minus pre-treatment evaluations. The results closely resemble the main ANCOVA estimates. Both AI-based and hybrid safety systems reduce perceived job satisfaction, meaning, and social value relative to the human-supervisor condition. The main qualitative conclusions, therefore, do not depend on the choice of estimation strategy. The only notable difference concerns perceived fair pay, where the hybrid treatment generates a larger negative coefficient than in the main specification.

Table 4: Robustness check: specifications with changes of the outcome variables

	(1) ΔJob satisfaction	(2) ΔMeaning	(3) ΔSocial value	(4) ΔFair wage	(5) ΔJob satisfaction	(6) ΔMeaning	(7) ΔSocial value	(8) ΔFair wage
AI vs. Human	-0.147***	-0.144***	-0.135**	-0.064	-0.153***	-0.149***	-0.146***	-0.065
	(0.053)	(0.053)	(0.053)	(0.053)	(0.054)	(0.053)	(0.053)	(0.053)
Hybrid vs. Human	-0.119**	-0.138***	-0.148***	-0.109**	-0.124**	-0.140***	-0.153***	-0.109**
	(0.050)	(0.051)	(0.051)	(0.053)	(0.051)	(0.051)	(0.052)	(0.054)
Controls	No	No	No	No	Yes	Yes	Yes	Yes
Observations	2,172	2,172	2,172	2,172	2,172	2,172	2,172	2,172
R-squared	0.004	0.004	0.004	0.002	0.013	0.017	0.012	0.011

Notes: All dependent variables are standardized to have a mean of 0 and a standard deviation of 1. Dependent variables are post-treatment minus pre-treatment changes. Controls include age, gender, and indicators for: higher education degree, married, working status, urban residence, living alone, children in the household, and high-income tercile. Robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Third, I examine whether respondents' engagement with the questionnaire affects the results. After completing the survey, the LISS panel asked respondents whether they found the questionnaire interesting and whether they enjoyed completing it. Both questions used a five-point scale ranging from 1 ("Not at all") to 5 ("Very much"). I include both variables as controls because they capture respondents' engagement with the survey and help account for systematic response styles, such as satisficing, acquiescence, or a tendency to rely on the endpoints of the response scales. Because five respondents did not answer these questions, I assigned them to a separate missing category to preserve the full estimation sample.

The results presented in Table 5 are highly similar to those in Models (5)-(8) of Table 2. Adding the survey evaluation controls changes the treatment coefficients only marginally. The estimated effects of hybrid safety supervision remain virtually identical across all outcomes. The estimated effects of AI-only safety supervision also remain similar, although the coefficient for perceived social value decreases slightly in magnitude and its statistical significance falls from the 1% to the 5% level. Overall, the findings indicate that respondents' engagement with the questionnaire and their response styles do not drive the main results.

Table 5: Robustness check: survey evaluation controls

	(1) Job satisfaction	(2) Meaning	(3) Social value	(4) Fair wage
AI vs. Human	-0.156***	-0.111***	-0.094**	-0.017
	(0.045)	(0.041)	(0.037)	(0.022)
Hybrid vs. Human	-0.124***	-0.112***	-0.100***	-0.043*
	(0.043)	(0.039)	(0.035)	(0.022)
Interest in the survey topic, omitted category 1 = not at all interesting				
2	0.025	0.161	0.026	0.016
	(0.112)	(0.119)	(0.101)	(0.052)
3	0.022	0.236*	0.138	-0.020
	(0.115)	(0.122)	(0.103)	(0.052)
4	0.094	0.224*	0.069	0.000
	(0.119)	(0.126)	(0.107)	(0.051)
Very interesting	-0.029	0.217	0.020	-0.020
	(0.130)	(0.139)	(0.119)	(0.053)
Missing interesting	0.372**	-0.299	-0.187	0.126
	(0.182)	(0.489)	(0.467)	(0.175)
Pleasure in filling out survey, omitted category 1 = not at all pleasurable				
2	0.041	-0.064	0.034	-0.040
	(0.141)	(0.126)	(0.108)	(0.076)
3	0.126	-0.067	0.027	0.024
	(0.140)	(0.134)	(0.112)	(0.068)
4	0.192	-0.016	0.120	0.001
	(0.145)	(0.140)	(0.117)	(0.069)
Very pleasurable	0.328**	0.066	0.227*	0.048
	(0.154)	(0.151)	(0.129)	(0.070)
Controls	Yes	Yes	Yes	Yes
Baseline outcome	Yes	Yes	Yes	Yes
Observations	2,172	2,172	2,172	2,172
R-squared	0.296	0.429	0.536	0.828

Notes: All dependent variables are standardized to have a mean of 0 and a standard deviation of 1. Controls include age, gender, and indicators for: higher education degree, married, working status, urban residence, living alone, children in the household, and high-income tercile. The evaluations of a survey control include a category for missing observations, which is not estimable in the case of how pleasurable respondents found filling out the survey. Robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

As an additional robustness check, I control for the Big Five personality traits. These variables capture stable individual differences that may influence how respondents evaluate hypothetical jobs and respond to survey questions. Table 6 shows that adding the personality measures leaves the estimated treatment effects virtually unchanged. Compared with Models (5) to (8) in Table 2, the treatment coefficients change very little. The only notable difference is that the estimated effect of hybrid safety supervision on perceived fair wages becomes slightly larger in magnitude and reaches the 5% significance level. Overall, the results indicate that stable personality differences do not explain the main findings.

Table 6: Robustness check: personality controls

	(1) Job satisfaction	(2) Meaning	(3) Social value	(4) Fair wage
AI vs. Human	-0.165***	-0.119***	-0.104***	-0.022
	(0.046)	(0.042)	(0.038)	(0.022)
Hybrid vs. Human	-0.125***	-0.110***	-0.100***	-0.049**
	(0.044)	(0.040)	(0.036)	(0.022)
Extraversion	0.022	0.005	-0.010	-0.004
	(0.021)	(0.019)	(0.017)	(0.010)
Agreeableness	0.088***	0.070***	0.072***	0.003
	(0.023)	(0.021)	(0.019)	(0.010)
Conscientiousness	0.025	0.005	0.020	-0.019*
	(0.023)	(0.020)	(0.018)	(0.010)
Emotional stability	0.029	0.018	0.008	0.000
	(0.023)	(0.019)	(0.018)	(0.009)
Imagination/Intellect	-0.056**	-0.039*	-0.038*	-0.001
	(0.024)	(0.022)	(0.020)	(0.010)
Controls	Yes	Yes	Yes	Yes
Baseline outcome	Yes	Yes	Yes	Yes
Observations	2,097	2,097	2,097	2,097
R-squared	0.308	0.434	0.543	0.832

Notes: All dependent variables and personality traits are standardized to have a mean of 0 and a standard deviation of 1. Controls include age, gender, and indicators for: higher education degree, married, working status, urban residence, living alone, children in the household, and high-income tercile. Robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Taken together, the robustness checks in Tables 3-6 show that the main findings remain stable. Restricting the sample to “attentive” respondents, estimating change-score models, and controlling for survey engagement and personality traits leave the estimated treatment effects largely unchanged. Across all specifications, AI-only and hybrid safety systems consistently reduce perceived job satisfaction, work meaningfulness, and social value relative to human supervision. The evidence for perceived fair wages remains weak, although the hybrid treatment yields a small negative effect across several robustness checks.

7. Heterogeneity Analyses

Tables 7-10 examine whether the treatment effects differ by respondents' demographic characteristics and personality traits. Overall, the results provide little evidence of systematic treatment effect heterogeneity. The joint F-tests generally fail to reject the null hypothesis that the treatment effects do not vary by age, gender, employment status, education, household income, or the Big Five personality traits.

One exception concerns job satisfaction, where the treatment effects differ by household income (p = 0.005). In particular, respondents in the highest-income tercile report a larger reduction in perceived job satisfaction under the hybrid safety system than respondents in the lower-income groups. A second exception concerns perceived fair wages, in which the joint interaction test for imagination/intellect is statistically significant (p = 0.026). Respondents scoring higher on imagination/intellect assign slightly higher fair wages to jobs with AI-based and hybrid safety systems than respondents scoring lower on this trait. However, given the large number of interaction tests, some statistically significant results are expected by chance alone. Overall, the evidence suggests that the effects of AI-based and hybrid safety systems remain broadly similar across demographic groups and personality traits.

Table 7: Heterogeneity by age and gender

	Moderator = Age above 58				Moderator = Female			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Job satisfaction	Meaning	Social value	Fair wage	Job satisfaction	Meaning	Social value	Fair wage
AI vs. Human	-0.187***	-0.108**	-0.148***	-0.037	-0.156**	-0.153***	-0.103**	0.008
	(0.062)	(0.054)	(0.046)	(0.034)	(0.064)	(0.058)	(0.050)	(0.030)
Hybrid vs. Human	-0.073	-0.064	-0.087*	-0.034	-0.165***	-0.193***	-0.138***	-0.031
	(0.060)	(0.054)	(0.048)	(0.034)	(0.063)	(0.056)	(0.049)	(0.030)
Moderator	0.097	0.037	-0.079	-0.005	0.030	-0.002	0.039	0.011
	(0.089)	(0.081)	(0.072)	(0.041)	(0.063)	(0.057)	(0.052)	(0.032)
AI x Moderator	0.047	-0.020	0.083	0.040	-0.013	0.070	-0.003	-0.049
	(0.091)	(0.081)	(0.073)	(0.044)	(0.091)	(0.081)	(0.073)	(0.044)
Hybrid x Moderator	-0.108	-0.101	-0.035	-0.015	0.076	0.156**	0.066	-0.019
	(0.086)	(0.078)	(0.071)	(0.043)	(0.086)	(0.078)	(0.071)	(0.044)
F-test p-value: interactions	0.198	0.391	0.240	0.412	0.556	0.136	0.544	0.536
Observations	2,172	2,172	2,172	2,172	2,172	2,172	2,172	2,172
R-squared	0.294	0.432	0.537	0.830	0.293	0.432	0.537	0.830

Notes: All dependent variables are standardized to have a mean of 0 and a standard deviation of 1. All models include the corresponding baseline outcome and the full set of controls, which are omitted from the table for brevity. Controls include age, gender, and indicators for: higher education degree, married, working status, urban residence, living alone, children in the household, and high-income tercile. Robust standard errors in parentheses. Reported F-test p-values correspond to a joint test that the coefficients on the AI × moderator and Hybrid × moderator interaction terms are both equal to zero. The median age in the sample is 58, which is why we use that age cutoff for the moderator. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 8: Heterogeneity by working status and income

	Moderator = Working				Moderator = Top-tertile income			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Job satisfaction	Meaning	Social value	Fair wage	Job satisfaction	Meaning	Social value	Fair wage
AI vs. Human	-0.129*	-0.107*	-0.071	0.028	-0.123**	-0.111**	-0.077*	-0.011
	(0.067)	(0.059)	(0.056)	(0.030)	(0.055)	(0.051)	(0.046)	(0.028)
Hybrid vs. Human	-0.093	-0.106*	-0.079	-0.015	-0.033	-0.089*	-0.059	-0.055**
	(0.063)	(0.056)	(0.052)	(0.028)	(0.052)	(0.049)	(0.045)	(0.026)
Moderator	0.053	-0.034	0.057	0.036	0.075	-0.013	0.007	0.001
	(0.069)	(0.063)	(0.057)	(0.036)	(0.076)	(0.069)	(0.063)	(0.040)
AI x Moderator	-0.070	-0.022	-0.071	-0.091**	-0.120	-0.017	-0.085	-0.019
	(0.090)	(0.081)	(0.072)	(0.044)	(0.095)	(0.084)	(0.075)	(0.045)
Hybrid x Moderator	-0.066	-0.013	-0.049	-0.054	-0.298***	-0.076	-0.140*	0.046
	(0.086)	(0.078)	(0.070)	(0.044)	(0.092)	(0.082)	(0.073)	(0.049)
F-test p-value: interactions	0.663	0.964	0.601	0.117	0.005	0.634	0.157	0.357
Observations	2,172	2,172	2,172	2,172	2,172	2,172	2,172	2,172
R-squared	0.293	0.431	0.537	0.830	0.296	0.432	0.537	0.830

Notes: All dependent variables are standardized to have a mean of 0 and a standard deviation of 1. All models include the corresponding baseline outcome and the full set of controls, which are omitted from the table for brevity. Controls include age, gender, and indicators for: higher education degree, married, working status, urban residence, living alone, children in the household, and high-income tercile. Robust standard errors in parentheses. Reported F-test p-values correspond to a joint test that the coefficients on the AI × moderator and Hybrid × moderator interaction terms are both equal to zero. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 9: Heterogeneity by education, extraversion, and agreeableness

	Moderator = Higher education				Moderator = Extraversion				Moderator = Agreeableness			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
	Job satisfaction	Meaning	Social value	Fair wage	Job satisfaction	Meaning	Social value	Fair wage	Job satisfaction	Meaning	Social value	Fair wage
AI vs. Human	-0.244***	-0.174***	-0.145***	0.009	-0.176***	-0.128***	-0.115***	-0.020	-0.170***	-0.123***	-0.110***	-0.021
	(0.061)	(0.056)	(0.050)	(0.031)	(0.046)	(0.041)	(0.037)	(0.022)	(0.046)	(0.041)	(0.037)	(0.022)
Hybrid vs. Human	-0.153***	-0.178***	-0.138***	-0.037	-0.132***	-0.116***	-0.108***	-0.049**	-0.128***	-0.112***	-0.105***	-0.049**
	(0.059)	(0.055)	(0.049)	(0.030)	(0.044)	(0.040)	(0.036)	(0.022)	(0.043)	(0.040)	(0.036)	(0.022)
Moderator	-0.090	-0.117**	-0.064	-0.008	0.066**	0.058**	0.052**	0.002	0.115***	0.118***	0.096***	-0.016
	(0.065)	(0.060)	(0.053)	(0.032)	(0.029)	(0.027)	(0.022)	(0.013)	(0.037)	(0.031)	(0.029)	(0.019)
AI x Moderator	0.176*	0.127	0.082	-0.053	-0.059	-0.063	-0.078**	-0.036*	-0.038	-0.093**	-0.062	-0.003
	(0.091)	(0.082)	(0.073)	(0.044)	(0.048)	(0.045)	(0.038)	(0.021)	(0.050)	(0.045)	(0.040)	(0.025)
Hybrid x Moderator	0.055	0.147*	0.066	-0.005	-0.011	-0.051	-0.061*	0.010	-0.044	-0.068	-0.035	0.035
	(0.087)	(0.078)	(0.071)	(0.044)	(0.044)	(0.039)	(0.034)	(0.021)	(0.048)	(0.042)	(0.039)	(0.025)
F-test p-value: interactions	0.149	0.132	0.483	0.397	0.451	0.272	0.066	0.105	0.631	0.092	0.292	0.177
Observations	2,159	2,159	2,159	2,159	2,097	2,097	2,097	2,097	2,097	2,097	2,097	2,097
R-squared	0.292	0.427	0.536	0.829	0.301	0.431	0.539	0.832	0.305	0.434	0.542	0.832

Notes: All dependent variables and personality traits are standardized to have a mean of 0 and a standard deviation of 1. All models include the corresponding baseline outcome and the full set of controls, which are omitted from the table for brevity. Controls include age, gender, and indicators for: higher education degree, married, working status, urban residence, living alone, children in the household, and high-income tercile. Reported F-test p-values correspond to a joint test that the coefficients on the AI × moderator and Hybrid × moderator interaction terms are both equal to zero. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 10: Heterogeneity by conscientiousness, emotional stability, and imagination/intellect

	Moderator = Conscientiousness				Moderator = Emotional stability				Moderator = Imagination/Intellect			
	(1) Job satisfaction	(2) Meaning	(3) Social value	(4) Fair wage	(5) Job satisfaction	(6) Meaning	(7) Social value	(8) Fair wage	(9) Job satisfaction	(10) Meaning	(11) Social value	(12) Fair wage
AI vs. Human	-0.168***	-0.123***	-0.110***	-0.022	-0.176***	-0.127***	-0.114***	-0.020	-0.174***	-0.126***	-0.114***	-0.018
	(0.047)	(0.042)	(0.037)	(0.022)	(0.046)	(0.041)	(0.037)	(0.022)	(0.046)	(0.042)	(0.037)	(0.022)
Hybrid vs. Human	-0.126***	-0.112***	-0.105***	-0.049**	-0.133***	-0.117***	-0.108***	-0.049**	-0.130***	-0.113***	-0.107***	-0.047**
	(0.044)	(0.040)	(0.036)	(0.023)	(0.044)	(0.040)	(0.036)	(0.022)	(0.044)	(0.040)	(0.036)	(0.022)
Moderator	0.062*	0.033	0.024	-0.010	0.076**	0.046	0.026	0.001	0.018	0.004	0.027	-0.041**
	(0.036)	(0.036)	(0.034)	(0.019)	(0.033)	(0.031)	(0.029)	(0.017)	(0.034)	(0.033)	(0.027)	(0.018)
AI x Moderator	-0.013	-0.025	0.004	-0.023	-0.060	-0.050	-0.030	-0.026	-0.037	-0.004	-0.047	0.042*
	(0.048)	(0.047)	(0.041)	(0.025)	(0.048)	(0.043)	(0.041)	(0.022)	(0.050)	(0.046)	(0.040)	(0.023)
Hybrid x Moderator	-0.041	-0.015	0.018	-0.005	-0.041	-0.020	-0.009	0.009	-0.025	-0.028	-0.053	0.059***
	(0.048)	(0.044)	(0.041)	(0.025)	(0.047)	(0.043)	(0.039)	(0.024)	(0.047)	(0.044)	(0.038)	(0.022)
	-0.168***	-0.123***	-0.110***	-0.022	-0.176***	-0.127***	-0.114***	-0.020	-0.174***	-0.126***	-0.114***	-0.018
F-test p-value: interactions	0.684	0.862	0.882	0.590	0.436	0.501	0.753	0.248	0.747	0.790	0.330	0.026
Observations	2,097	2,097	2,097	2,097	2,097	2,097	2,097	2,097	2,097	2,097	2,097	2,097
R-squared	0.300	0.430	0.539	0.832	0.300	0.430	0.539	0.832	0.298	0.430	0.539	0.832

Notes: All dependent variables and personality traits are standardized to have a mean of 0 and a standard deviation of 1. All models include the corresponding baseline outcome and the full set of controls, which are omitted from the table for brevity. Controls include age, gender, and indicators for: higher education degree, married, working status, urban residence, living alone, children in the household, and high-income tercile. Robust standard errors in parentheses. Reported F-test p-values correspond to a joint test that the coefficients on the AI × moderator and Hybrid × moderator interaction terms are both equal to zero.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

8. Additional Outcomes: Dignity and Privacy

As a post hoc analysis, I also examine whether respondents believe that the new safety system respects workers' dignity and privacy and whether they consider it effective. Table 11 shows large, highly statistically significant effects on dignity and privacy. Compared with human supervision, AI-only safety systems reduce perceived respect for workers' dignity by about 0.40 standard deviations and for privacy by about 0.45 standard deviations. The hybrid safety system produces nearly identical effects. These effect sizes are approximately three to four times larger than those for job satisfaction, work meaningfulness, and perceived social value. By contrast, respondents perceive the effectiveness of AI-based and hybrid safety systems as broadly comparable to that of human supervision. Together, these findings suggest that respondents' reservations about AI-based workplace safety systems stem primarily from concerns about dignity and privacy rather than doubts about their effectiveness. The fact that participants do not perceive a difference in the effectiveness of AI- and human-based systems likely explains why they do not think a higher wage is necessary to compensate workers supervised by AI.

Table 11: Additional post-treatment outcomes: dignity, privacy, and effectiveness in improving workplace safety

	(1) Dignity	(2) Privacy	(3) Effectiveness	(4) Dignity	(5) Privacy	(6) Effectiveness
AI vs. Human	-0.394***	-0.439***	-0.081	-0.404***	-0.452***	-0.100*
	(0.052)	(0.051)	(0.054)	(0.052)	(0.051)	(0.054)
Hybrid vs. Human	-0.366***	-0.431***	0.031	-0.377***	-0.442***	0.031
	(0.052)	(0.051)	(0.052)	(0.051)	(0.050)	(0.052)
Controls	No	No	No	Yes	Yes	Yes
Observations	2,167	2,167	2,167	2,167	2,167	2,167
R-squared	0.032	0.042	0.002	0.047	0.063	0.034

Notes: All dependent variables are standardized to have a mean of 0 and a standard deviation of 1. All models include the corresponding baseline outcome and the full set of controls, which are omitted from the table for brevity. Controls include age, gender, and indicators for: higher education degree, married, working status, urban residence, living alone, children in the household, and high-income tercile. Robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

9. Discussion and Conclusion

AI increasingly supports or replaces human supervision in workplace health and safety, yet little evidence exists on how individuals perceive these alternative systems and how they may influence the perceived quality of work. This paper provides the first causal evidence on this question using a pre-registered vignette experiment embedded in a nationally representative Dutch survey. Respondents evaluated otherwise identical jobs that differed only in whether workplace safety relied on a human supervisor, an AI-based system, or a hybrid AI-human system.

The results show that respondents consistently evaluate AI-based and hybrid safety systems less favorably than human supervision. Relative to human supervision, both AI-only and hybrid systems reduce perceived job satisfaction, work meaningfulness, and the perceived social value of the job. Contrary to the expectation that retaining a human supervisor (i.e., a “human-in-the-loop”) would mitigate these effects, respondents evaluate the hybrid system almost identically to the AI-only system. These findings suggest that introducing AI into workplace safety changes how respondents evaluate work quality, even when humans remain involved in the supervisory process.

The post hoc analyses provide additional insight into these perceptions. Respondents perceive AI-based and hybrid safety systems as substantially less respectful of workers' dignity and privacy than human supervision. These effects are approximately three to four times as large as the corresponding effects on job satisfaction, meaningfulness, and social value. By contrast, respondents do not perceive AI-based systems as substantially more or less effective than human supervision. Taken together, these findings suggest that respondents' concerns focus primarily on the implications of AI for workers' dignity and privacy rather than on the technology's ability to improve workplace safety.

The evidence for compensating wage differentials is considerably weaker. Respondents do not systematically assign higher fair wages to jobs that rely on AI-based workplace safety systems, even though they evaluate these jobs as less meaningful and less satisfying. This finding contrasts with the standard prediction from compensating wage differential theory and remains difficult to explain. One possibility is that respondents separate evaluations of job quality from normative judgments about appropriate wages. Another possibility is that respondents view wages primarily as compensation for productivity rather than for non-pecuniary job characteristics. The present study cannot distinguish between these explanations, and future research should examine why reductions in perceived job quality do not translate into higher perceived fair wages.

More broadly, the findings contribute to the emerging literature on the non-pecuniary consequences of AI in the workplace. Much of the existing debate focuses on productivity, employment, and wages. The results presented here suggest that AI also shapes how individuals evaluate the human experience of work. Even when respondents consider AI-based safety systems to be as effective as human supervision, they perceive them as reducing important non-monetary aspects of job quality. These perceptions may reflect broader concerns about opaque algorithmic decision-making, reduced human interaction, and diminished respect for workers' autonomy, privacy, and dignity.

Several limitations suggest promising directions for future research. First, respondents evaluated hypothetical job scenarios rather than actual workplace experiences. Although the experimental design provides strong internal validity, future studies should examine whether similar effects emerge after workers gain direct experience with AI-based safety systems. Second, the study focuses on the Netherlands, a country characterized by relatively safe workplaces, strong labor-market institutions, and comparatively high levels of institutional trust. Whether these findings generalize to countries with different institutional settings remains an open question. Third, respondents evaluated a single hypothetical job profile. Future research could employ paired-vignette or conjoint designs (Hainmuller et al., 2015) to increase external validity and study the trade-offs individuals make when evaluating alternative workplace technologies. Finally, future work should examine whether familiarity with AI-based safety systems changes these perceptions over time or whether concerns about privacy, dignity, and work quality persist as these technologies become more common.

Overall, the findings suggest that organizations introducing AI into workplace safety should pay attention not only to the technology's effectiveness but also to how workers perceive its implications for work quality. Designing AI systems that preserve workers' dignity, privacy, and human agency may prove as important for successful adoption as improving the systems' technical performance.

Datasets used in this paper

Centerdata. (2026). *Wave 17: LISS Core Study Personality (Project 7.17).* LISS panel, Tilburg University. https://doi.org/10.57990/m3s9-vx93

Centerdata. (2025). Impact of AI-driven Safety Systems on Worker Well-being (L_WorkerSafetySystems_1.0p). LISS panel, Tilburg University. https://www.dataarchive.lissdata.nl/study-units/view/1720

Centerdata. (2024). *LISS Core Study: Personality, Wave 16 (7.16).* LISS panel, Tilburg University, Centerdata. https://doi.org/10.57990/c63p-1250

Lloyd's Register Foundation. (2024). World Risk Poll 2023. https://www.lrfoundation.org.uk/wrp/world-risk-poll-data.

APPENDIX A
Survey Questionnaire *Impact of AI-driven Safety Systems on Worker Well-being*

DUTCH VERSION OF THE QUESTIONNAIRE (ENGLISH BELOW)

Introductie (alle respondenten)
In deze vragenlijst leest u over een fabriek en de veiligheidsregels die daar gelden. Stel u zich voor dat u als werknemer in deze fabriek werkt. Beantwoord de vragen op basis daarvan zo goed mogelijk. Uw antwoorden helpen ons beter te begrijpen hoe werknemers veiligheid op de werkvloer ervaren.

Standaardbeschrijving (alle respondenten)
Een 40-jarige man werkt 40 uur per week aan de lopende band in een grote fabriek. Hij is getrouwd en verkeert in goede lichamelijke en geestelijke gezondheid. Zijn taken bestaan uit het bevestigen van onderdelen, het controleren op defecten of het inpakken van producten terwijl deze over een transportband bewegen.

Het bedrijf houdt zich aan alle veiligheidsregels uit de Arbowet. Werknemers dragen beschermende uitrusting (handschoenen, helmen, gehoorbescherming) en krijgen veiligheidstrainingen. De machines worden goed onderhouden. De werksfeer is positief en collegiaal. De taken bestaan vaak uit het tillen van zware voorwerpen, langdurig staan en het werken in lawaai van machines. Soms zijn de productieschema's strak en moet er snel worden gewerkt om deadlines te halen.

Basisvragen (alle respondenten)

Lees de bovenstaande beschrijving zorgvuldig door. Als u het antwoord niet zeker weet, lees dan de beschrijving opnieuw voordat u verdergaat.

1. Volgens de beschrijving, welke veiligheidsregels zijn aanwezig voor werknemers in de fabriek?

- **<u>Werknemers dragen beschermende uitrusting, krijgen veiligheidstraining en de machines worden onderhouden</u>**
- Werknemers nemen hun eigen gereedschap mee en bepalen hun eigen veiligheidsregels
- Werknemers krijgen geen veiligheidstraining en inspecties zijn zeldzaam
- Weet ik niet

2. Op basis van de beschrijving, in welke mate is het werk van deze werknemer…?
(1 = Helemaal niet, 7 = Heel erg)

	1	2	3	4	5	6	7
…veilig	○	○	○	○	○	○	○
…risicovol	○	○	○	○	○	○	○

	1 2 3 4 5 6 7
…stressvol	○ ○ ○ ○ ○ ○ ○
…fysiek zwaar	○ ○ ○ ○ ○ ○ ○

3. Op basis van de beschrijving, hoe waarschijnlijk is het dat dit werk voor deze werknemer leidt tot …
(1 = Zeer onwaarschijnlijk, 4 = Niet onwaarschijnlijk en niet waarschijnlijk, 7 = Zeer waarschijnlijk)

	1 2 3 4 5 6 7
…arbeidsongevallen	○ ○ ○ ○ ○ ○ ○
…lichamelijke klachten (bijv. rugpijn, nekpijn, oogklachten)	○ ○ ○ ○ ○ ○ ○
…mentale klachten (bijv. stress, vermoeidheid, angst)	○ ○ ○ ○ ○ ○ ○

4. Op basis van de beschrijving, in welke mate denkt u dat de werknemer…
(1 = Helemaal niet, 7 = Heel erg veel)

	1 2 3 4 5 6 7
tevreden is met het werk	○ ○ ○ ○ ○ ○ ○
het werk zinvol vindt	○ ○ ○ ○ ○ ○ ○
het werk maatschappelijk nuttig vindt	○ ○ ○ ○ ○ ○ ○

5. Op basis van de beschrijving, wat vindt u een passend **bruto maandsalaris** voor deze werknemer? *Ga uit van een 40-urige werkweek en een salaris exclusief bonussen.*

- ○ Minder dan €1.500
- ○ €1.500–1.999
- ○ €2.000–2.499
- ○ €2.500–2.999
- ○ €3.000–3.499
- ○ €3.500–3.999
- ○ €4.000 of meer

De fabriek voert een nieuw veiligheidssysteem in. Op de volgende pagina krijgt u een beschrijving van dit nieuwe systeem. Lees de beschrijving goed door en beantwoord dan de vragen over dit nieuwe systeem.
Beschrijving nieuw veiligheidssysteem

Groep 1: Alleen AI-veiligheidssysteem (N = 700)
Het bedrijf voert een nieuw veiligheidssysteem in dat volledig is gebaseerd op kunstmatige intelligentie (AI), speciaal ontwikkeld om in real-time risico's en gevaren op de werkvloer te detecteren.

Het AI-systeem houdt de productievloer voortdurend in de gaten met behulp van camera's en slimme wearables (zoals helmen met sensoren), en monitort veiligheid gerelateerde signalen zoals vermoeidheid, oververhitting of plotselinge, onregelmatige bewegingen. Het controleert ook op gevaren zoals defecte machines, gladde vloeren of het ontbreken van beschermende uitrusting.

Het AI-systeem geeft directe veiligheidswaarschuwingen en aanbevelingen aan werknemers (bijv. het nemen van een pauze). Het systeem en alle verzamelde gegevens worden uitsluitend gebruikt voor veiligheidsdoeleinden.

Groep 2: Alleen menselijk veiligheidssysteem (N = 700)

Het bedrijf voert een nieuw veiligheidssysteem in dat volledig is gebaseerd op getrainde veiligheidstoezichthouders, speciaal ontwikkeld om in real-time risico's en gevaren op de werkvloer te detecteren.

De toezichthouders houden de productievloer voortdurend in de gaten door regelmatig rondes te lopen en te letten op aan veiligheid gerelateerde signalen zoals vermoeidheid, oververhitting of plotselinge, onregelmatige bewegingen. Ze controleren ook op gevaren zoals defecte machines, gladde vloeren of het ontbreken van beschermende uitrusting.

De getrainde toezichthouders geven directe veiligheidswaarschuwingen en aanbevelingen aan werknemers (bijv. het nemen van een pauze). Het systeem en alle verzamelde gegevens worden uitsluitend gebruikt voor veiligheidsdoeleinden.

Groep 3: Hybride veiligheidssysteem (AI + mens) (N = 700)

Het bedrijf voert een nieuw veiligheidssysteem in dat een combinatie is van kunstmatige intelligentie (AI) en getrainde veiligheidstoezichthouders, speciaal ontwikkeld om in real-time risico's en gevaren op de werkvloer te detecteren.

Het AI-systeem houdt de productievloer voortdurend in de gaten met behulp van camera's en slimme wearables (zoals helmen met sensoren), en monitort veiligheid gerelateerde signalen zoals vermoeidheid, oververhitting of plotselinge, onregelmatige bewegingen. Het controleert ook op gevaren zoals defecte machines, gladde vloeren of het ontbreken van beschermende uitrusting.

De toezichthouders houden de productievloer voortdurend in de gaten door regelmatig rondes te lopen en te letten op aan dezelfde veiligheid gerelateerde signalen en gevaren. Zij beoordelen ook de waarschuwingen van het AI-systeem en geven directe veiligheidswaarschuwingen en aanbevelingen aan werknemers (bijv. het nemen van een pauze). Het systeem en alle verzamelde gegevens worden uitsluitend gebruikt voor veiligheidsdoeleinden.

Groepsspecifieke vraag

Lees de bovenstaande beschrijving zorgvuldig door. Als u het antwoord niet zeker weet, lees dan de beschrijving opnieuw voordat u verdergaat.

6. Volgens de beschrijving, waaruit bestaat het nieuwe veiligheidssysteem?

6a. [Alleen getoond aan Groep 1: Alleen AI-veiligheidssysteem]

○ Een op AI gebaseerd veiligheidssysteem
○ Meer gebruik van veiligheidsbrillen en -schoenen
○ Nieuwe veiligheidstraining
○ Weet ik niet

6b. [Alleen getoond aan Groep 2: Alleen menselijk veiligheidssysteem]

○ Getrainde veiligheidstoezichthouders
○ Meer gebruik van veiligheidsbrillen en -schoenen
○ Nieuwe veiligheidstraining
○ Weet ik niet

6c. [Alleen getoond aan Groep 3: Hybride veiligheidssysteem (AI + mens)]

○ Een combinatie van AI en getrainde veiligheidstoezichthouders
○ Meer gebruik van veiligheidsbrillen en -schoenen
○ Nieuwe veiligheidstraining
○ Weet ik niet

Vragen voor alle respondenten

7. Na de invoering van het nieuwe veiligheidssysteem, in welke mate is het werk van deze werknemer…

(1 = Helemaal niet, 7 = Zeer veel)

	1	**2**	**3**	**4**	**5**	**6**	**7**
…veilig	○	○	○	○	○	○	○
…risicovol	○	○	○	○	○	○	○
…stressvol	○	○	○	○	○	○	○
…fysiek zwaar	○	○	○	○	○	○	○

8. Na de invoering van het nieuwe veiligheidssysteem, hoe waarschijnlijk is het dat dit werk voor deze werknemer leidt tot…

(1 = Zeer onwaarschijnlijk, 4 = Niet onwaarschijnlijk en niet waarschijnlijk, 7 = Zeer waarschijnlijk)

	1 2 3 4 5 6 7
…arbeidsongevallen	○ ○ ○ ○ ○ ○ ○
…lichamelijke klachten (bijv. nekpijn, rugpijn, oogklachten)	○ ○ ○ ○ ○ ○ ○
…mentale klachten (bijv. stress, vermoeidheid, angst)	○ ○ ○ ○ ○ ○ ○

9. Na de invoering van het nieuwe veiligheidssysteem, in welke mate denkt u dat de werknemer…

(1 = Helemaal niet, 7 = Zeer veel)

	1 2 3 4 5 6 7
…tevreden is met het werk	○ ○ ○ ○ ○ ○ ○
…het werk zinvol vindt	○ ○ ○ ○ ○ ○ ○
…het werk maatschappelijk nuttig vindt	○ ○ ○ ○ ○ ○ ○

10. Na invoering van het nieuwe veiligheidssysteem, wat vindt u een passend **bruto maandsalaris** voor deze werknemer? *Ga uit van een 40-urige werkweek en een salaris exclusief bonussen.*

- ○ Minder dan €1.500
- ○ €1.500–1.999
- ○ €2.000–2.499
- ○ €2.500–2.999
- ○ €3.000–3.499
- ○ €3.500–3.999
- ○ €4.000 of meer

11. In welke mate hebt u er vertrouwen in dat het nieuwe veiligheidssysteem…

(1 = Helemaal geen vertrouwen, 7 = Volledig vertrouwen)

	1 2 3 4 5 6 7
…arbeidsongevallen kan voorkomen	○ ○ ○ ○ ○ ○ ○
…mentale gezondheid op het werk kan verbeteren	○ ○ ○ ○ ○ ○ ○
…fysieke gezondheid op het werk kan verbeteren	○ ○ ○ ○ ○ ○ ○
…gerichte adviezen geven over veiligheid	○ ○ ○ ○ ○ ○ ○

12. Hoe effectief denkt u dat het nieuwe veiligheidssysteem is in het verbeteren van de veiligheid op de werkvloer?

(1 = Helemaal niet effectief, 7 = Heel erg effectief)

13. In welke mate respecteert het nieuwe veiligheidssysteem het volgende…

(1 = Helemaal niet, 7 = Volledig)

	1	**2**	**3**	**4**	**5**	**6**	**7**
…de privacy van werknemers	○	○	○	○	○	○	○
…de waardigheid van werknemers	○	○	○	○	○	○	○

ENGLISH QUESTIONNAIRE

Introduction (All respondents)
In this survey, you will read about a manufacturing plant and its safety measures. Imagine the experience of a worker at this plant. Please answer the questions based on this description as best as you can. Your answers will help us better understand how workers experience workplace safety.

Baseline Vignette (All respondents)
Imagine a 40-year-old male assembly line worker working 40 hours per week in a large manufacturing plant. He is married and in good physical and mental health. His tasks include attaching parts, checking for defects, or packing products as items move along a conveyor belt.

The company follows all workplace safety rules under the Dutch Occupational Health and Safety Act (Arbowet).

Workers wear protective gear (gloves, helmets, ear protection) and receive safety training. The machines are well maintained. The workplace culture is positive and cooperative. The tasks often include carrying heavy loads, standing for long periods, and noise from the machines. Production schedules can be tight, and the worker sometimes has to work quickly to meet deadlines.

Baseline Questions (All respondents)
Please read the description above carefully. If you are not sure of the answer, re-read the description before responding.

1. According to the description, what safety measures are in place for workers in the manufacturing plant?
   <u>Workers wear protective gear, receive safety training, and machinery is maintained</u>
   Workers bring their own tools and decide their own safety rules
   Workers do not receive safety training, and inspections are rare
   I don't know

2. Based on the description, to what extent do you think that this job is…?

	1= Not at all	**2**	**3**	**4**	**5**	**6**	**7 = Very much**
Safe	o	o	o	o	o	o	o
Risky	o	o	o	o	o	o	o
Stressful	o	o	o	o	o	o	o
Physically demanding	o	o	o	o	o	o	o

3. Based on the description, how likely do you think it is that this job leads to the following for this worker?

	1 = Very Unlikely	2	3	4 = Neither Likely nor Unlikely	5	6	7 = Very Likely
Workplace accidents	o	o	o	o	o	o	o
Physical health issues (e.g., neck pain, back pain, eye strain)	o	o	o	o	o	o	o
Mental health issues (e.g., stress, fatigue, anxiety)	o	o	o	o	o	o	o

4. Based on the description, to what extent do you think that this worker …?

	1 = Not at all	2	3	4	5	6	7 = Very much
Is satisfied with the job	o	o	o	o	o	o	o
Finds the job meaningful	o	o	o	o	o	o	o
Finds the job socially useful	o	o	o	o	o	o	o

5. Based on the description, what do you think is an appropriate <u>monthly</u> gross salary for this worker? *Assume a 40-hour workweek and a salary excluding bonuses.*
   Less than €1,500
   €1,500–1,999
   €2,000–2,499
   €2,500–2,999
   €3,000–3,499
   €3,500–3,999
   €4,000 or more

The manufacturing plant introduces a new safety system. On the next page you will see a description of the new system. Read the description carefully and answer the questions about the new system.

Description of the new safety system

3. Treatment Scenarios (Random Assignment to 3 Groups, roughly 700 respondents per group)

Group 1: AI-Only Safety System, N=700
The company introduces a new safety system based entirely on Artificial Intelligence (AI), specifically designed to detect workplace risks and hazards in real time.

The AI system continuously monitors the factory floor using smart wearables (e.g., helmets with sensors) and cameras. It tracks safety-related signs such as fatigue, overheating, or sudden irregular movements, and checks for hazards such as malfunctioning machinery, slippery floors, or missing safety gear.

The AI system provides immediate safety alerts and recommendations to workers (e.g., suggesting a break). The safety system and any collected data are used only for safety purposes.

Group 2: Human-Only Safety System, N=700
The company introduces a new safety system based entirely on trained safety supervisors, specifically designed to detect workplace risks and hazards in real time.

The supervisors continuously monitor the factory floor by regularly walking through it and observing safety-related signs such as fatigue, overheating, or sudden irregular movements, and checking for hazards such as malfunctioning machinery, slippery floors, or missing safety gear.

The trained safety supervisors provide immediate safety alerts and recommendations to workers (e.g., suggesting a break). The safety system and any collected data are used only for safety purposes.

Group 3: Hybrid Human + AI System, N=700
The company introduces a new safety system based on a combination of Artificial Intelligence (AI) and trained safety supervisors, specifically designed to detect workplace risks and hazards in real time.

The AI system continuously monitors the factory floor using smart wearables (e.g., helmets with sensors) and cameras. It tracks safety-related signs such as fatigue, overheating, or sudden irregular movements, and checks for hazards such as malfunctioning machinery, slippery floors, or missing safety gear.

The trained safety supervisors regularly walk through the factory floor to observe workers directly and look for the same types of hazards. They also review AI safety alerts and provide immediate recommendations to workers (e.g., suggesting a break). The safety system and any collected data are used only for safety purposes.

Group-specific question
Please read the description above carefully. If you are not sure of the answer, re-read the description before responding.

6. According to the description, what does the new safety system consist of?
   6a**. [Shown only to Group 1: AI-Only Safety System]**
      <u>An AI-based safety system</u>
      Increased use of safety goggles and boots
      New safety training
      I don't know
   6b**. [Shown only to Group 2: Human-Only Safety System]**
      <u>Trained safety supervisors</u>
      Increased use of safety goggles and boots
      New safety training
      I don't know

   6c. **[Shown only to Group 3: Hybrid Human + AI System]**
      <u>A combination of AI and trained safety supervisors</u>
      Increased use of safety goggles and boots
      New safety training
      I don't know

Questions for all respondents

Answer all questions based on the description you just read.

7. After the implementation of the new safety system, to what extent do you think that this job is…?

	1= Not at all	2	3	4	5	6	7 = Very much
Safe	o	o	o	o	o	o	o
Risky	o	o	o	o	o	o	o
Stressful	o	o	o	o	o	o	o
Physically demanding	o	o	o	o	o	o	o

8. After the implementation of the new safety system, how likely do you think it is that this job leads to the following for this worker…?

	1 = Very Unlikely	2	3	4 = Neither Likely nor Unlikely	5	6	7 = Very Likely
Workplace accidents	o	o	o	o	o	o	o
Physical health issues (e.g., neck pain, back pain, eye strain)	o	o	o	o	o	o	o

	1 = Very Unlikely	2 3	4 = Neither Likely nor Unlikely	5 6	7 = Very Likely
Mental health issues (e.g., stress, fatigue, anxiety)	o	o o o		o o o	

9. After the implementation of the new safety system, to what extent to you think that this worker …?

	1=Not at all	2	3	4	5	6	7 = Very
Is satisfied with the job	o	o	o	o	o	o	o
Finds the job meaningful	o	o	o	o	o	o	o
Finds the job socially useful	o	o	o	o	o	o	o

10. After the implementation of the new safety system, what do you think is an appropriate <u>monthly</u> gross salary for this worker? *Assume a 40-hour workweek and a salary excluding bonuses.*
    Less than €1,500
    €1,500–1,999
    €2,000–2,499
    €2,500–2,999
    €3,000–3,499
    €3,500–3,999
    €4,000 or more

11. How much would you trust the new safety system to do the following…?

	1= Completely distrust	2 3 4 5 6	7 = Completely trust
Prevent workplace accidents	o	o o o o o o	
Improve mental health at work	o	o o o o o o	
Improve physical health at work	o	o o o o o o	
Provide targeted safety recommendations	o	o o o o o o	

12. How effective do you think the new safety system is in improving workplace safety?
    1 = Not Effective – 7 = Very Effective

13. Based on the description, to what extent do you think this new safety system respects the following?

	1 = Not at all	**2**	**3**	**4**	**5**	**6**	**7 = Completely**
Workers' privacy	○	○	○	○	○	○	○
Workers' dignity	○	○	○	○	○	○	○

References

Abeliansky, A. L., Beulmann, M., & Prettner, K. (2024). Are they coming for us? Industrial robots and the mental health of workers. *Research Policy*, *53*(3), 104956. https://doi.org/https://doi.org/10.1016/j.respol.2024.104956

Akerlof, G. A., & Yellen, J. L. (1990). The Fair Wage-Effort Hypothesis and Unemployment. *The Quarterly Journal of Economics*, *105*(2), 255–283. https://doi.org/10.2307/2937787

Ariely, D., Kamenica, E., & Prelec, D. (2008). Man's search for meaning: The case of Legos. *Journal of Economic Behavior & Organization*, *67*(3-4), 671–677.

Berente, N., Gu, B., Recker, J., & Santhanam, R. (2021). Managing artificial intelligence. *MIS quarterly*, *45*(3), 1433–1450.

Bick, A., Blandin, A., Deming, D. J., Fuchs-Schündeln, N., & Jessen, J. (2026). *Mind the Gap: AI Adoption in Europe and the US*.

Böckerman, P., & Ilmakunnas, P. (2009). Job Disamenities, Job Satisfaction, Quit Intentions, and Actual Separations: Putting the Pieces Together. *Industrial Relations: A Journal of Economy and Society*, *48*(1), 73–96. https://doi.org/10.1111/j.1468-232X.2008.00546.x

Bryson, A., Kauhanen, A., & Rouvinen, P. (2026). AI and Worker Well-being: Evidence from a Nationally Representative Study.

Burbano, V., Padilla, N., & Meier, S. (2024). Gender differences in preferences for meaning at work. *American Economic Journal: Economic Policy*, *16*(3), 61–94.

Burbano, V. C. (2016). Social responsibility messages and worker wage requirements: Field experimental evidence from online labor marketplaces. *Organization Science*, *27*(4), 1010–1028.

Cassar, L., & Meier, S. (2018). Nonmonetary Incentives and the Implications of Work as a Source of Meaning. *Journal of Economic Perspectives*, *32*(3), 215–238.

Clark, A., Georgellis, Y., & Sanfey, P. (2012). Job Satisfaction, Wage Changes, and Quits: Evidence from Germany. In S. W. Polachek & K. Tatsiramos (Eds.), *35th Anniversary Retrospective* (Vol. 35, pp. 499–525). Emerald Group Publishing Limited. https://doi.org/10.1108/S0147-9121(2012)0000035041

Clark, A. E. (2001). What really matters in a job? Hedonic measurement using quit data. *Labour economics*, *8*(2), 223–242. https://doi.org/https://doi.org/10.1016/S0927-5371(01)00031-8

Cnossen, F., & Nikolova, M. (2025). Work meaningfulness and effort. *Journal of Behavioral and Experimental Economics*, *119*, 102460. https://doi.org/https://doi.org/10.1016/j.socec.2025.102460

Cornelißen, T. (2009). The Interaction of Job Satisfaction, Job Search, and Job Changes. An Empirical Investigation with German Panel Data [Article]. *Journal of happiness studies*, *10*(3), 367–384. https://doi.org/10.1007/s10902-008-9094-5

Das, M., & Knoef, M. (2019). Experimental and longitudinal data for scientific and policy research: open access to data collected in the longitudinal Internet Studies for the Social sciences (LISS) panel. *Data-Driven Policy Impact Evaluation: How Access to Microdata is Transforming Policy Design*, 131–146.

De Neve, J.-E., Diener, E., Tay, L., & Xuereb, C. (2013). The objective benefits of subjective well-being. In J. F. Helliwell, R. Layard, & J. Sachs (Eds.), *World Happiness Report* (pp. 54–79).

De Simone, M. (2026). Automation and Health and Safety of Workers (HSW): Evidence from Literature Review. *Economia & lavoro*, *60*(1), 81–101.

De Simone, M., Guarascio, D., & Reljic, J. (2025). The impact of robots on workplace injuries and deaths. Empirical evidence from Europe. *LEM WORKING PAPER SERIES*, 1–58.

Dixon, J., Hong, B., & Wu, L. (2021). The robot revolution: Managerial and employment consequences for firms. *Management Science*, *67*(9), 5586–5605.

Dur, R., & Van Lent, M. (2019). Socially useless jobs. *Industrial Relations: A Journal of Economy and Society*, *58*(1), 3–16.

El-Helaly, M. (2024). Artificial intelligence and occupational health and safety, benefits and drawbacks. *La Medicina del lavoro*, *115*(2), e2024014.

Gihleb, R., Giuntella, O., Stella, L., & Wang, T. (2022). Industrial robots, Workers' safety, and health. *Labour economics*, *78*, 102205. https://doi.org/https://doi.org/10.1016/j.labeco.2022.102205

Giuntella, O., Konig, J., & Stella, L. (2025). Artificial intelligence and the wellbeing of workers. *Scientific Reports*, *15*(1), 20087.

Goldberg, L. R. (1992). The development of markers for the Big-Five factor structure. *Psychological Assessment*, *4*(1), 26.

Gunadi, C., & Ryu, H. (2021). Does the rise of robotic technology make people healthier? *Health economics*, *30*(9), 2047–2062. https://doi.org/https://doi.org/10.1002/hec.4361

Hainmueller, J., Hangartner, D., & Yamamoto, T. (2015). Validating vignette and conjoint survey experiments against real-world behavior. *Proceedings of the National Academy of Sciences*, *112*(8), 2395–2400.

Hendriks, M., & Cnossen, F. (2026). The economics of meaningful work: A scoping review. In M. Nikolova (Ed.), *Work Meaning and Motivation: Interdisciplinary Perspectives on Work Well-being* (pp. 151–184). Springer.

Jetha, A., Bakhtari, H., Irvin, E., Biswas, A., Smith, M. J., Mustard, C., Arrandale, V. H., Dennerlein, J. T., & Smith, P. M. (2025). Do occupational health and safety tools that utilize artificial intelligence have a measurable impact on worker injury or illness? Findings from a systematic review. *Systematic reviews*, *14*(1), 146.

Kesternich, I., Schumacher, H., Siflinger, B., & Schwarz, S. (2021). Money or meaning? Labor supply responses to work meaning of employed and unemployed individuals. *European Economic Review*, *137*, 103786. https://doi.org/https://doi.org/10.1016/j.euroecorev.2021.103786

Küsters, A., & Schneider, B. (2025). What is technological unemployment. *University of Oxford Working Paper*. https://doi.org/https://ora.ox.ac.uk/objects/uuid:839abc6e-1c7a-48f2-aff6-0fad792b44f9

La Torre, G., Manai, M. V., Meucci, S., Lucente, A., Picerno, A., Ammirati, S., & De Sio, S. (2026). Artificial intelligence and occupational health and safety: a systematic review. *Journal of Public Health*, 1–21.

Lévy-Garboua, L., Montmarquette, C., & Simonnet, V. (2007). Job satisfaction and quits. *Labour economics*, *14*(2), 251–268. https://doi.org/https://doi.org/10.1016/j.labeco.2005.08.003

LRF. (2024). *World Risk Poll 2024 report: Engineering safer workplaces: Global trends in occupational safety and health*. https://wrp.lrfoundation.org.uk

Mokyr, J., Vickers, C., & Ziebarth, N. L. (2015). The History of Technological Anxiety and the Future of Economic Growth: Is This Time Different? *Journal of Economic Perspectives*, *29*(3), 31–50. https://doi.org/10.1257/jep.29.3.31

Nikolova, M., & Cnossen, F. (2020). What makes work meaningful and why economists should care about it. *Labour economics*, *65*, 101847. https://doi.org/https://doi.org/10.1016/j.labeco.2020.101847

Nikolova, M., Cnossen, F., & Nikolaev, B. (2024). Robots, meaning, and self-determination. *Research Policy*, *53*(5), 104987. https://doi.org/https://doi.org/10.1016/j.respol.2024.104987

Nikolova, M., & de Wit, J. (2026). Work orientations and turnover. *Journal of Population Economics*, *39*(3), 33. https://doi.org/10.1007/s00148-026-01176-w

Nikolova, M., Lepinteur, A., & Cnossen, F. (2025). Just another cog in the machine? A worker-level view of robotization and tasks. *Economica*. https://doi.org/10.1111/ecca.70006

Nikolova, M., Milanova, V., & Wang, F. (2026). When AI Does the Work: Does Attribution Shape Meaning and Effort? *Working Paper*.

Nikolova, M., Nikolaev, B., & Boudreaux, C. (2023). Being your own boss and bossing others: the moderating effect of managing others on work meaning and autonomy for the self-employed and employees. *Small Business Economics*, *60*(2), 463–483. https://doi.org/10.1007/s11187-021-00597-z

Non, A., Rohde, I., de Grip, A., & Dohmen, T. (2022). Mission of the company, prosocial attitudes and job preferences: A discrete choice experiment. *Labour economics*, *74*, 102087. https://doi.org/https://doi.org/10.1016/j.labeco.2021.102087

Noponen, N., Feshchenko, P., Auvinen, T., Luoma-aho, V., & Abrahamsson, P. (2024). Taylorism on steroids or enabling autonomy? A systematic review of algorithmic management. *Management Review Quarterly*, *74*(3), 1695–1721. https://doi.org/10.1007/s11301-023-00345-5

Probst, T., Shoss, M., Bankins, S., Behrend, T., Dragano, N., Jetha, A., Koopman, J., Selenko, E., Tang, P. M., & Tay, L. (2026). Aligning artificial intelligence with worker health, well-being, and safety: An occupational health psychology agenda. *Occupational Health Science*, *10*(2), 15.

Riso, S., Adăscăliței, D., & Oosterwijk, G. (2025). *Implications of algorithmic management for work, employment and social dialogue: Literature review*.

https://www.eurofound.europa.eu/en/publications/all/implications-of-algorithmic-management-for-work-employment-and-social-dialogue-literature-review

Rosen, S. (1986). Chapter 12 The theory of equalizing differences. In O. C. Ashenfelter & R. Layard (Eds.), *Handbook of Labor Economics* (Vol. 1, pp. 641–692). Elsevier. https://doi.org/https://doi.org/10.1016/S1573-4463(86)01015-5

Rosso, B. D., Dekas, K. H., & Wrzesniewski, A. E. (2010). On the meaning of work: A theoretical integration and review. *Research in organizational behavior*, *30*, 91–127.

Sadeghian, S., & Hassenzahl, M. (2022). The" artificial" colleague: evaluation of work satisfaction in collaboration with non-human coworkers. proceedings of the 27th international conference on intelligent user interfaces,

Scherpenzeel, A. (2011). Data collection in a probability-based internet panel: how the LISS panel was built and how it can be used. *Bulletin of Sociological Methodology/Bulletin de Méthodologie Sociologique*, *109*(1), 56–61.

Schouwer, T. D., Gsottbauer, E., Kesternich, I., & Schumacher, H. (2025). Work meaning and fair wages. *Labour economics*, *97*, 102808. https://doi.org/https://doi.org/10.1016/j.labeco.2025.102808

Schwabe, H., & Castellacci, F. (2020). Automation, workers' skills and job satisfaction. *PloS one*, *15*(11), e0242929.

Soffia, M., Leiva-Granados, R., Zhou, X., & Skordis, J. (2024). *From technology exposure to job quality: Evidence from a comprehensive UK survey*. https://www.ifow.org/publications/from-technology-exposure-to-job-quality-evidence-from-a-comprehensive-uk-survey

Stantcheva, S. (2023). How to run surveys: A guide to creating your own identifying variation and revealing the invisible. *Annual Review of Economics*, *15*(1), 205–234.

Toole, H., Hutton, E., Zhang, K., Wong, C., Smith, P., Maistrello, G., Fahy, N., & Nightingale, M. (2025). *The intersection of occupational safety and health with emerging technology: Scoping and mapping the evidence to plan for future research*. https://www.lrfoundation.org.uk/sites/default/files/2025-08/emerging-tech-technical.pdf

Zhang, L., & Amos, C. (2026). Artificial intelligence and work meaningfulness. *Behaviour & Information Technology*, 1–15.