CHAPTER 10

AI-Mediated Performance Appraisal and Employee Perceptions of Fairness: Evidence from Pune's IT Sector

Bhakti Pardeshi and Meeta Meshram***

ABSTRACT

Integrating artificial intelligence (AI) into how organizations assess job performance is getting a lot of attention from both academics and people in the field. A key question is whether workers think it's fair, since their careers could be affected by what these AI systems decide. This research looks at how employees in the information technology (IT) industry in Pune, India, feel about the fairness of AI-driven performance reviews, both in terms of the process (how it works) and the outcomes (what results it produces). We used a survey with specific questions to gather data from a group of 301 IT workers in Pune. The survey included questions about procedural justice, asking about things like transparency, whether employees had a say, consistency, how well things were explained, and if there was a way to appeal decisions. It also included questions about distributive justice, focusing on whether the outcomes, rewards, promotions, and recognition were fair, and how accurate the evaluations seemed to be. For each question, we looked at how often people chose each answer. Then, to analyze the data, we used t-tests to compare the average scores for each aspect of fairness (procedural and distributive) to a neutral score of 3.00. The results showed that, on average, employees felt the procedures used in AI-based appraisals were fairer than neutral, suggesting that they generally accepted some parts of how these systems work. However, when it came to distributive justice, opinions were more varied. While many agreed with some aspects, there were still worries about whether the outcomes were fair and how accurate the evaluations were. These findings suggest that just being transparent and explaining how AI systems work isn't enough. Organizations also need to have strong ways for employees to address concerns and regularly check the outcomes to make sure things are fair. This study adds local evidence from the IT sector in India and suggests practical steps that companies can take to make AI in HR more accepted.

Keywords: AI-based performance appraisal; Procedural justice; Distributive justice; Algorithmic fairness; IT sector in Pune.

1.0 Introduction

The quick spread of AI tech within people work power management has changed how work checks are done.

**Corresponding author; Student, Department of HR, Neville Wadia Institute of Management and Research, Pune, Maharashtra, India (E-mail: bhaktip24@gmail.com)*

***Professor, Department of HR, Neville Wadia Institute of Management and Research, Pune, Maharashtra, India (E-mail: meetamee9@gmail.com)*

DOI: 10.17492/JPI/ICISI2025/251210 *ISBN: 978-93-49790-69-8*

Old school checks were led by bosses at set times, but now, they use AI tools that offer more speed, fairness, and reach. In the IT world—known for its heavy use of data, digital marks, and high need for work and new ideas—AI-based check systems are very much liked. These systems use data from code houses, issue lists, work sites, and work tools to make work scores, suggest marks, and even hint at ways to get better. While these steps may cut some human mistakes and work load, they also bring up new worries about being clear, easy to read, and fair in how people feel.

How fair things seem is key to how firms do. Justice in how things are done, results are given, and how people are treated helps shape work drive, trust in bosses, job happy feel, and plans to stay at the work place. When checks affect rises, money extras, and work paths, seeing unfairness can kill work drive and cause people to leave. Adding AI to these checks makes it more complex by using hard codes, auto weights, and data ways that folks may not get or agree with. Some think that AI systems may repeat old data mistakes or store boss likes in hard-to-see ways, making it hard to say they're safer.

Though there is more study on AI in work control and watching, we still don't always know how AI-based check systems change how fair workers think things are, mainly in IT parts in certain places. Pune, India—with many IT firms from huge global ones to local new ones—is a good place to look into this. Firms here use digital work marks a lot different and let computers make check choices in many different ways. Yet, past research has mostly missed how workers in these new market places really feel about AI-influenced check ways, where local views, normal ways, and work markets may shape ideas of fairness differently than in western places.

This study plans to dig into how AI-driven check ways and workers' ideas of fairness link in Pune's IT world. It has two main goals: first, to map out how AI bits are fit into current check ways; and second, to look into how different types of fairness are seen by workers dealing with these systems. Using ideas from how firms see justice and how social-tech systems work, the research guesses that seeing things clearer and having workers help design checks will link to better views of process fairness, while hard-to-see AI ways and few ways to fix wrongs will link to bad views on results and treatment fairness.

By method, the study mixes surveys and talks to truly get what workers feel. This way, we can count clear links while still getting the full picture needed to make sense of AI-touched talks. By focusing on Pune's IT folk, the research adds local proof that adds to world chats about good AI and fair HR numbers.

The need for this study is both from thinking and doing. By thoughts, it adds to what we know about org justice by looking into how digital go-betweens change key justice ideas at tech-heavy work places. By action, it gives HR folks and org bosses clear advice on design steps—like being open steps, ways to involve workers, and ways to appeal—that

DOI: 10.17492/JPI/ICISI2025/251210 ISBN: 978-93-49790-69-8

may make bad views less and make AI-check systems seem more right. By doing so, the study helps guide talks on rules for AI choices in work places, saying that using tech well must go with fair steps to keep org trust and work results. In short, as firms use more AI to handle work better, knowing how these tech steps affect people becomes key. By looking into how AI-driven checks impact fairness ideas in a major Indian IT place, this study hopes to give deep and useful insight, making AI use in worker practices fairer.

2.0 Review of Literature

Beigang (2022) looks into how we split fairness in predicting and giving out roles through AI, saying mixing these two makes us miss key choices. For AI in job ratings, this is big: these systems guess work quality, but their results decide raises, jobs, or firings. Beigang's work thus lets us ask if workers care more about accurate guesses or fair results, and it leads to ways to test these separately in studies.

Corbett-Davies and Goel (2018) review common fairness ways in AI, noting many standards don't fit well together or miss social and tech harms. They warn that one fix won't capture "fairness," and we must pick tools based on the situation and values. Their ideas back up using many fairness types and mix methods linking stats and worker views to find gaps in fairness and what's accepted.

Dastin (2018) shares a big case where an AI in hiring was unfair to women, causing it to stop. This case in the review shows how choices in data and AI can keep old biases, shape trust and needs care in tech and talks when setting up rating AI in places like Pune's tech scene. Kleinberg, Ludwig, Mullainathan, and Rambachan (2018) look at deep issues in AI fairness, discussing unavoidable choices and effects. This is vital for job rating studies as it pushes looking into what fairness workers value and if job rules really show these choices. Lee (2018) studies how we see AI decisions linked to trust, emotions, and actions. For job rating AI studies, Lee shows how seen fairness, or not having a say, might change feelings and trust in a firm, affecting job happiness or staying at a job.

Park, Ahn, Hosanagar, and Lee (2021) dive into how we react to AI in jobs and why workers might push back, pointing out fairness and trust issues. Their findings work well in Pune, suggesting clearer AI roles, more user say, and ongoing checks could ease concerns and boost fairness—points this study will check in action.

Park, Ahn, Hosanagar, and Lee (2022) build on past work by offering ways to design fair AI in jobs, looking at goals and shared choices. This shapes the ethics part of the review, giving ideas for steps (like challenges, shared design, clear rules) whose effects on how workers see fairness can be studied. Patel *et al.* (2022) talk about an AI system for ranking worker performance, showing how tech setups and choices lead to ranked results.

DOI: 10.17492/JPI/ICISI2025/251210 ISBN: 978-93-49790-69-8

Their tech details help in the review by discussing what firms use and how it feels fair or marks “lower” workers. Qin *et al.* (2023) compare how we see human bosses and AI in rating jobs, giving new data on trust and fairness among types of rating. Their new findings guide guesses about AI’s pluses and downs in making fair, clear, and balanced job calls among workers. Robert, Pierce, Marquis, Kim, and Alahmad (2020) cover criticism of AI in managing staff and suggest study and design paths for better systems. Their review is key to the thesis as it links deep ideas and real steps, giving a base to measure firm rules in Pune’s firms. Schlicker *et al.* (2021) test how showing AI workings changes fairness views, showing that clear explanations can switch how fair AI seems, but limits exist. Their results urge using detailed surveys, noting clear rules alone may not win good fairness views.

Selbst *et al.* (2019) call for focusing on real office practices and power in fairness studies, not just abstract ideas. Their angle supports a mix of tech checks and real worker stories to grasp how fairness plays out and is fought over in real settings.

Shin (2020) offers tested ways to measure how we see AI choices on fairness, clear rules, and more. For current research, Shin’s tools give tried ways to track how Pune workers rate AI job rating systems. Starke, Baleis, Keller, and Marcinkowski (2022) sum up studies on how we see AI in making choices, noting key things that shape views. Their review adds to this study by setting real tests for parts like clear rules and making fair calls in rising IT markets. Whittlestone, Nyrup, Alexandrova, Dihal, and Cave (2019) lay out the big picture on AI ethics and suggest a research plan focusing on rules, public talks, and mixing fields. Their policy view adds to the thesis, linking study results on worker fairness views to wider rule suggestions for fair HR technology.

Together, these reviews show three linked ideas that help us learn about AI in job reviews and how fair people think they are: first, fairness in tech HR work has many sides and often faces hard choices, like choosing between types of fairness or dealing with different stats rules. So, studies need to split tech quality from how it is spread and not just see “fairness” as one simple idea; second, how workers see fairness isn’t just about if the model is right but also about parts like being clear, easy to understand, chances to speak up and fix issues, and how companies use these tools—this means checking should mix well-tested survey parts (like how fair the process is, results are, how they treat you; how clear they are) with deep talks about real life and who holds power; and third, fixing problems needs both tech and rules design: making things clear and open to challenge is key but not enough, it must also fit with shared rule-making, checks, and safe rules to keep trust. All of these points suggest a research plan that uses both numbers and stories, one that clearly tells apart guessing right from effects, tests how changes (like more details, ways to object, making together) shift how fair people think it is, and turns what we learn into tips for companies and rules that fit places like Pune’s tech area.

DOI: 10.17492/JPI/ICISI2025/251210 ISBN: 978-93-49790-69-8

3.0 Objectives and Hypotheses

3.1 Objectives

- To examine employees' perceptions of procedural fairness in AI-mediated performance appraisal systems in Pune's IT sector.
- To evaluate employees' perceptions of distributive fairness arising from AI-mediated performance appraisal systems in Pune's IT sector.

3.2 Hypotheses

- H1: The mean score for procedural justice (composite of five Likert items) among IT employees is significantly different from the neutral midpoint (3.0).
- H2: The mean score for distributive justice (composite of five Likert items) among IT employees is significantly different from the neutral midpoint (3.0).

4.0 Research Methodology

A simple, clear survey type was used. We gathered basic data with set questions given to chosen IT workers (n = 301) in Pune. The form asked staff to rate things from 1 (Strongly Disagree) to 5 (Strongly Agree). Five questions checked on fair processes and five on fair outcomes. We made easy-to-read charts for each answer. To test deeper, we used one-sample t-tests on average scores (means of five questions each) and compared them to a mid-point of 3.00. We handled data and did the math using common tools; we did not do any linking or predicting tests.

5.0 Data Analysis

Table 1: Procedural Item 1 (Transparency)

Response	1 (Strongly Disagree)	2 (Disagree)	3 (Neutral)	4 (Agree)	5 (Strongly Agree)
Frequency	12	20	48	110	111
Percent (%)	3.99	6.64	15.95	36.54	36.88

Mean = 3.957; SD = 1.075.

Explanation: The item got an average score of 3.957 (SD = 1.075). The percent break-up shows that 36.54% of the people said yes, and 36.88% said a big yes. On the other side, 3.99% were in firm no, and 6.64% just said no. These numbers show that most people had good thoughts about this item; the top picks were Agree/Strongly Agree. It looks like, in general, people felt good about this part of the AI-guided check, but a small group was still unsure or didn't like it.

DOI: 10.17492/JPI/ICISI2025/251210 ISBN: 978-93-49790-69-8

Table 2: Procedural Item 2 (Voice)

Response	1 (Strongly Disagree)	2 (Disagree)	3 (Neutral)	4 (Agree)	5 (Strongly Agree)
Frequency	14	26	55	100	106
Percent (%)	4.65	8.64	18.27	33.22	35.22

Mean = 3.857; SD = 1.133.

Meaning: The score averaged at 3.857 (SD = 1.133). Percent break down shows 33.22% agreed and 35.22% agreed a lot, but 4.65% did not agree at all and 8.64% just disagreed. These numbers show that most people had good thoughts about this; the top picks were Agree/Strongly Agree. The trend tells us that people generally like this part of AI-helped checks, even though a small group sat on the fence or didn't like it.

Table 3: Procedural Item 3 (Consistency)

Response	1 (Strongly Disagree)	2 (Disagree)	3 (Neutral)	4 (Agree)	5 (Strongly Agree)
Frequency	10	22	60	105	104
Percent (%)	3.32	7.31	19.93	34.88	34.55

Mean = 3.9; SD = 1.063.

Meaning: The score was an average of 3.9 (SD = 1.063). The breakdown of numbers shows 34.88% of people said they agreed and 34.55% said they agreed a lot. On the other side, 3.32% did not agree at all and 7.31% just did not agree. Most people thought well of this part; the top picks were Agree or Strongly Agree. The spread hints that most folks liked this part of the AI-guided check, but a small group was still unsure or did not like it.

Table 4: Procedural Item 4 (Explanation)

Response	1 (Strongly Disagree)	2 (Disagree)	3 (Neutral)	4 (Agree)	5 (Strongly Agree)
Frequency	16	30	70	95	90
Percent (%)	5.32	9.97	23.26	31.56	29.9

Mean = 3.708; SD = 1.152.

Meaning: The score was 3.708 with a spread of 1.152. The data shows 31.56% of people said yes and 29.9% said a strong yes. On the other hand, 5.32% did not like it at all and 9.97% just said no. More people feel good than bad about this point; most said they agree or strongly agree. The numbers show most people see this part of AI-based review in a good light, but some are still on the fence or don't like it.

DOI: 10.17492/JPI/ICISI2025/251210 ISBN: 978-93-49790-69-8

Table 5: Procedural Item 5 (Right to appeal)

Response	**1 (Strongly Disagree)**	**2 (Disagree)**	**3 (Neutral)**	**4 (Agree)**	**5 (Strongly Agree)**
Frequency	18	28	66	89	100
Percent (%)	5.98	9.3	21.93	29.57	33.22

Mean = 3.748; SD = 1.184.

Meaning: The item got an average score of 3.748 (SD = 1.184). The percent split shows about 29.57% agreed and 33.22% strongly agreed. On the other side, 5.98% strongly disagreed and 9.3% just disagreed. These numbers show that most people felt good about this item; the top picks were Agree/Strongly Agree. So, it looks like most people liked this part of AI-looking stuff, even though some were unsure or didn't like it.

Table 6: Distributive Item 1 (Outcomes)

Response	**1 (Strongly Disagree)**	**2 (Disagree)**	**3 (Neutral)**	**4 (Agree)**	**5 (Strongly Agree)**
Frequency	20	30	70	95	86
Percent (%)	6.64	9.97	23.26	31.56	28.57

Mean = 3.654; SD = 1.183.

Explanation: The item got an average score of 3.654 (SD = 1.183). The data shows that 31.56% of people agreed and 28.57% agreed very much. On the other hand, 6.64% did not agree at all and 9.97% just disagreed. These numbers show that most people had good thoughts about this item; the top picks were Agree/Strongly Agree. The spread of scores hints that most people like this part of AI-based review. But still, a good chunk of people were either in-between or did not like it.

Table 7: Distributive Item 2 (Rewards)

Response	**1 (Strongly Disagree)**	**2 (Disagree)**	**3 (Neutral)**	**4 (Agree)**	**5 (Strongly Agree)**
Frequency	22	34	80	88	77
Percent (%)	7.31	11.3	26.58	29.24	25.58

Mean = 3.545; SD = 1.195.

Reading: The item got an average score of 3.545 (SD = 1.195). The percents show that 29.24% of people said yes and 25.58% said a big yes, while 7.31% said a big no and

11.3% said no. These numbers show that most of the people felt good about this item; the top picks were Agree/Strongly Agree. The layout hints that people mostly liked this part of the AI-assisted review, though a small group was either neutral or did not like it.

Table 8: Distributive Item 3 (Promotions)

Response	1 (Strongly Disagree)	2 (Disagree)	3 (Neutral)	4 (Agree)	5 (Strongly Agree)
Frequency	24	26	68	92	91
Percent (%)	7.97	8.64	22.59	30.56	30.23

Mean = 3.664; SD = 1.218.

Understanding: The average score for this item was 3.664, with a standard gap of 1.218. The percent breakdown shows 30.56% said yes and 30.23% said a big yes, while 7.97% were totally against it and 8.64% said no. These numbers show that most people had a good view of this item; the top choices were Agree/Strongly Agree. The way these numbers spread out tells us that most people think well of this part of AI-based judgment, but a small group still felt so-so or not good.

Table 9: Distributive Item 4 (Recognition)

Response	1 (Strongly Disagree)	2 (Disagree)	3 (Neutral)	4 (Agree)	5 (Strongly Agree)
Frequency	15	25	75	105	81
Percent (%)	4.98	8.31	24.92	34.88	26.91

Mean = 3.704; SD = 1.103.

Meaning: The score for this item was on average 3.704 (with a standard deviation of 1.103). The percent break down shows that about 35% of people agreed and 27% agreed a lot. But, about 5% did not agree at all and 8% just disagreed. The numbers show that more people had good thoughts on this than bad ones; most picked Agree or Strongly Agree. The way these points lay out hints that most see this part of AI-led review in a good light, yet, a small but clear group still feels neutral or not happy about it.

Table 10: Distributive Item 5 (Accuracy)

Response	1 (Strongly Disagree)	2 (Disagree)	3 (Neutral)	4 (Agree)	5 (Strongly Agree)
Frequency	13	27	72	110	79
Percent (%)	4.32	8.97	23.92	36.54	26.25

Mean = 3.714; SD = 1.082.

DOI: 10.17492/JPI/ICISI2025/251210 ISBN: 978-93-49790-69-8

Here is what the data says: The average score was 3.714 with a spread of 1.082. About 36.54% agreed and 26.25% very much agreed with the item. On the flip side, 4.32% did not agree strongly and 8.97% just did not agree. Most of the replies were positive, with "Agree" and "Strongly Agree" being the most common picks. The data shows that most people felt good about this part of AI-help in judging, even though some were unsure or not happy.

4. One-Sample t-test (Consolidated per Hypothesis)

Table 11: Hypothesis 1 (Procedural Justice) Test Value = 3.00

t	df	Sig. (2-tailed)	Mean Difference	Std. Error of Mean	95% Confidence Interval of the Difference
28.823	300	0.0	0.834	0.0289	0.777 to 0.891

The mean is 3.834, with a rough SD of 0.502. We ran a one-sample t-test against the base value of 3.00. This gave $t(300) = 28.823$, with $p = 0.0$. This shows the mean score was well off from the middle mark of 3.00. The gap between them is 0.834.

Table 12: Hypothesis 2 (Distributive Justice) Test Value = 3.00

t	df	Sig. (2-tailed)	Mean Difference	Std. Error of Mean	95% Confidence Interval of the Difference
21.971	300	0.0	0.656	0.0299	0.597 to 0.715

Average score = 3.656 (about SD = 0.518). One tests of one sample vs. a middle value of 3.00 gave $t(300) = 21.971$, $p = 0.0$. This shows that the average score was very not like the middle. Average change = 0.656.

5.0 Findings

The study showed that, on the whole, IT workers in Pune think well of the step-by-step parts of AI-run job reviews. The average score for these steps was 3.834, and a test showed this score was way off the middle point ($t = 28.823$, $p = 0.0$). At the level of each item, more people agreed or strongly agreed than not, though some were in the middle or did not agree on things like being able to appeal. When it comes to fair sharing, the average score was 3.656. A test confirmed a big gap from the middle point here too ($t = 21.971$, $p = 0.0$). More people felt the results, rewards, and moves up were fair than not. But, views varied, especially on how right the reviews were, showing that some worries about fair share are still there. These facts suggest that while workers might be okay with some AI parts by step, doubts on the results stay and need rules to fix.

DOI: 10.17492/JPI/ICISI2025/251210 *ISBN: 978-93-49790-69-8*

6.0 Conclusion

This investigation explores how employees in Pune's IT industry view fairness when artificial intelligence is involved in performance evaluations. The study looked at two types of fairness: procedural (the fairness of the process) and distributive (the fairness of the outcomes). The results showed that employees generally thought the procedures were fair but had mixed feelings about whether the outcomes were distributed fairly.

These findings suggest that to maintain fairness in AI-driven evaluations, it's not enough to just focus on the technical aspects like making the AI transparent and understandable. Organizations also need to have systems in place that allow employees to appeal decisions and participate in how these systems are governed.

For those working in the field, this study suggests implementing clear explanations of how AI evaluations work, providing ways for employees to voice their opinions, and regularly checking the fairness of the evaluation outcomes. For researchers, the results suggest the need for more studies that combine different research methods to understand how cultural and organizational factors influence fairness perceptions. Future work should also explore how specific changes to the design of AI systems can impact these perceptions.

References

Beigang, F. (2022). On the advantages of distinguishing between predictive and allocative fairness in algorithmic decision-making. *Minds and Machines, 32*(4), 655–682. https://doi.org/10.1007/s11023-022-09615-9

Corbett-Davies, S., & Goel, S. (2018). The measure and mismeasure of fairness: A critical review of fair machine learning. *arXiv preprint* arXiv:1808.00023.

Dastin, J. (2018). Amazon scraps secret AI recruiting tool that showed bias against women. *Reuters*. Retrieved November 5, 2023, from https://www.reuters.com/article/world/insight-amazon-scraps-secret-ai-recruiting-tool-that-showed-bias-against-women-idUSKCN1MK0AG/

Kleinberg, J., Ludwig, J., Mullainathan, S., & Rambachan, A. (2018). Algorithmic fairness. *AEA Papers and Proceedings, 108*, 22–27. https://doi.org/10.1257/pandp.20181018

Lee, M. K. (2018). Understanding perception of algorithmic decisions: Fairness, trust, and emotion in response to algorithmic management. *Big Data & Society, 5*(1), 1–16. https://doi.org/10.1177/2053951718756684

Park, H., Ahn, D., Hosanagar, K., & Lee, J. (2021). Human–AI interaction in human resource management: Understanding why employees resist algorithmic evaluation at workplaces and how to mitigate burdens. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. https://doi.org/10.1145/3411764.3445304

Park, H., Ahn, D., Hosanagar, K., & Lee, J. (2022). Designing fair AI in human resource management: Understanding tensions surrounding algorithmic evaluation and envisioning stakeholder-centered solutions. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*. https://doi.org/10.1145/3491102.3517672

Patel, K., Sheth, K., Mehta, D., Tanwar, S., Florea, B. C., Taralunga, D. D., Altameem, A., Altameem, T., & Sharma, R. (2022). RanKer: An AI-based employee-performance classification scheme to rank and identify low performers. *Mathematics, 10*(19), Article 3714. https://doi.org/10.3390/math10193714

Qin, S., Jia, N., Luo, X., Liao, C., & Huang, Z. (2023). Perceived fairness of human managers compared with artificial intelligence in employee performance evaluation. *Journal of Management Information Systems, 40*(4), 1039–1070. https://doi.org/10.1080/07421222.2023.2267316

Robert, L. P., Pierce, C., Marquis, L., Kim, S., & Alahmad, R. (2020). Designing fair AI for managing employees in organizations: A review, critique, and design agenda. *Human–Computer Interaction, 35*(5–6), 545–575. https://doi.org/10.1080/07370024.2020.1735391

Schlicker, N., Langer, M., Ötting, S. K., Baum, K., König, C. J., & Wallach, D. (2021). What to expect from opening up 'black boxes'? Comparing perceptions of justice between human and automated agents. *Computers in Human Behavior, 122*, Article 106837. https://doi.org/10.1016/j.chb.2021.106837

Selbst, A. D., Boyd, D., Friedler, S. A., Venkatasubramanian, S., & Vertesi, J. (2019). Fairness and abstraction in sociotechnical systems. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*. https://doi.org/10.1145/3287560.3287598

Shin, D. (2020). User perceptions of algorithmic decisions in the personalized AI system: Perceptual evaluation of fairness, accountability, transparency, and explainability. *Journal of Broadcasting & Electronic Media, 64*(4), 541–565. https://doi.org/10.1080/08838151.2020.1843357

DOI: 10.17492/JPI/ICISI2025/251210 *ISBN: 978-93-49790-69-8*

Starke, C., Baleis, J., Keller, B., & Marcinkowski, F. (2022). Fairness perceptions of algorithmic decision-making: A systematic review of the empirical literature. *Big Data & Society, 9*(2), 1–16. https://doi.org/10.1177/20539517221115189

Whittlestone, J., Nyrup, R., Alexandrova, A., Dihal, K., & Cave, S. (2019). Ethical and societal implications of algorithms, data, and artificial intelligence: A roadmap for research. *Nuffield Foundation.*

DOI: 10.17492/JPI/ICISI2025/251210 *ISBN: 978-93-49790-69-8*