



A critique of the ethical effects of using AI in job recruitment and selection in South Africa.

Submitted by Sikhulile Angel Mzobe

Student Number: 217037405

Dissertation submitted in fulfilment of the requirements for the degree of

Master of Arts

In the

DISCIPLINE OF ETHICS

**School of Religion, Philosophy and Classics, University of KwaZulu-Natal,
Pietermaritzburg Campus, South Africa.**

November 2025

Supervised by


Prof. Beatrice Okyere-Manu

(Supervisor)


DECLARATION

I, Sikhulile Angel Mzobe, declare that:

- i. The research reported in this dissertation, except where otherwise indicated, is my original research.
- ii. This dissertation has not been submitted for any degree or examination at any other university.
- iii. This dissertation does not contain other persons' data, pictures, graphs, or other information, unless specifically acknowledged as being sourced from other persons.
- iv. This dissertation does not contain other persons' writing, unless specifically acknowledged as being sourced from other researchers. Where other written sources have been quoted:
 - Their words have been rewritten, but the general information attributed to them has been referenced.
 - Where their exact words have been used, their writing has been placed inside quotation marks and referenced.
- v. This dissertation does not contain text, graphics, or tables copied and pasted from the Internet, unless specifically acknowledged, and the source is detailed in the dissertation and in the References sections.

Signature: 

Date: 27 November 2025

Supervisor's signature: 

LIST OF ABBREVIATIONS

AI:	Artificial Intelligence
ATS:	Applicant Tracking System
B-BBEE:	Broad-Based Black Economic Empowerment
CEE:	Commission for Employment Equity
DECA:	Describe, Evaluate, Consult, Act
EEA:	Employment Equity Act (No. 55 of 1998)
HR:	Human Resources
PAJA:	Promotion of Administrative Justice Act (No. 3 of 2000)
POPIA:	Protection of Personal Information Act (No. 4 of 2013)
SA:	South Africa
SDG:	Sustainable Development Goal
UNESCO:	United Nations Educational, Scientific, and Cultural Organization
UNICEF:	United Nations International Children's Emergency Fund
WEF:	World Economic Forum

DEDICATION

May this work serve as a voice for those whose opportunities are shaped by AI recruitment systems, reminding us that technology must be guided by fairness, dignity, and justice. This study is offered in solidarity with your journey, and in hope that future innovations will open doors rather than close them.

ACKNOWLEDGEMENTS

I would like to express my deepest gratitude to everyone who supported me throughout the course of this dissertation.

First and foremost, I am profoundly thankful to my supervisor, Professor Beatrice Okyere-Manu, for her invaluable guidance, encouragement, and expertise. Her insights and unwavering support have been instrumental in shaping this research.

I extend heartfelt appreciation to my family, especially my father, Sibusiso Mzobe (Ngcolwane), whose steadfast emphasis on the importance of education laid the foundation for this pursuit. To my friends, thank you for your patience, love, and belief in me during the most challenging times. Your presence has been my anchor.

To my colleagues and peers, I am grateful for the stimulating discussions, constructive feedback, and camaraderie that enriched this academic journey.

This work is done in honour of two extraordinary women whose spirits walk with me: Mama Zamathembu Shezi, who never stopped believing in me even when almost everyone thought I would not even finish high school, and my late mother, Xolisile (Twana) Zulu, whose love, resilience, and unwavering belief in me laid the foundation for everything I am.

This dissertation is a tribute to their legacy, a quiet flame that burns in my heart and lights the path forward. Lastly, I acknowledge the institutions, authors, and scholars whose work laid the foundation for this study. Your contributions continue to inspire and inform.

ABSTRACT

The integration of Artificial Intelligence (AI) into recruitment and selection has transformed human resource management globally, introducing efficiency, precision, and scalability. In South Africa, this transformation occurs within a socio-economic and ethical context shaped by structural inequalities, high unemployment, and constitutional imperatives of fairness and redress. Traditional recruitment processes are often inefficient, with time-to-hire periods exceeding 90 days, and subjective evaluations create a compelling case for AI adoption.

This study critically examines the ethical implications of AI-driven recruitment in South Africa through a consequentialist lens, which evaluates actions based on their outcomes rather than their intentions. It explores whether AI enhances welfare, equity, and transparency or perpetuates bias, exclusion, and opacity.

Adopting a qualitative, exploratory, desktop-based design, the study uses secondary data from academic, policy, and institutional sources. Applying the DECA method (Describe, Evaluate, Consult, Act) and thematic analysis, it identifies four ethical dimensions: efficiency, fairness, transparency, and candidate dignity. Consequentialist reasoning evaluates AI's outcomes within South Africa's legal framework, encompassing the Employment Equity Act (EEA), the Protection of Personal Information Act (POPIA), and the Promotion of Administrative Justice Act (PAJA).

Findings reveal that while AI streamlines recruitment and reduces bias, it may also introduce algorithmic discrimination and diminish human oversight. The study concludes that AI's ethical legitimacy depends on advancing welfare, justice, and dignity. It recommends governance frameworks, capacity-building initiatives, and policies that ensure transparency, accountability, and human control in AI recruitment.

key words: *Artificial Intelligence, Recruitment, Ethics, Consequentialism, Fairness, Accountability*

TABLE OF CONTENTS

DECLARATION	i
LIST OF ABBREVIATIONS	iii
DEDICATION	iv
ACKNOWLEDGEMENTS	v
CHAPTER ONE: INTRODUCTION TO THE RESEARCH	11
1.0 Introduction.....	11
1.1 Background and Motivation	11
1.2 Research Problem	14
The ethical challenges and consequences that arise when Artificial Intelligence (AI) is used in recruitment and selection processes in South Africa.	14
1.3 Key Research Question.....	14
1.4 Research Sub-Questions	14
1.5 Key Research Objective.....	14
1.6 Research Sub-Objectives	14
1.7 Preview of Research Methodology and Method.....	14
1.8 Preview of Theoretical Framework: Theory of Consequentialism.....	15
1.9 Significance of the study.....	17
1.10 Definition of key concepts	20
1.10.1 Artificial Intelligence (AI)	20
1.10.2 Recruitment	20
1.10.3 Ethical Recruitment.....	20
1.10.4 Algorithmic Bias	21
1.10.5 Fairness.....	21
1.10.6 Transparency and Explainability.....	22
1.10.7 South Africa.....	22
1.11 Research Structure	23
1.12 Conclusion	26
CHAPTER TWO: LITERATURE REVIEW	27
2.0 Introduction.....	27
2.1 Traditional Recruitment Process.....	27
2.1.1 The Nature of the Traditional Recruitment Process.....	29
2.1.2 Stages of the Traditional Recruitment Process	31
2.1.3 Power of Traditional Recruitment Methods in South Africa	33

2.2	Limitations of the Traditional Recruitment Process	34
2.4.1	Overview of AI in Recruitment.....	42
2.4.2	The Nature of Artificial Intelligence in Recruitment	44
2.4.3	AI-Driven Recruitment Systems	45
2.4.4	Definition and Scope of Artificial Intelligence in Recruitment.....	46
2.4.5	Applications of Artificial Intelligence in Recruitment.....	47
2.5	Legal and Regulatory Compliance in Recruitment.....	49
2.5.1	South African Legislation on Employment and AI Recruitment.....	50
2.5.1.1	The Protection of Personal Information Act (POPIA) (Act No. 4 of 2013).....	51
2.5.1.2	The Employment Equity Act (EEA) (Act No. 55 of 1998).....	52
2.5.1.3	The Labour Relations Act (LRA) (Act No. 66 of 1995)	53
2.5.1.4	The Promotion of Administrative Justice Act (PAJA) (Act 3 of 2000)	54
2.5.2	Comparative World Ecstasy Monitor.....	55
2.5.2.1	Implication for South Africa.....	55
2.6	Gaps in Literature	56
2.7	Conclusion	57
CHAPTER THREE: METHOD AND METHODOLOGY		59
3.0	Introduction.....	59
3.1	Philosophical Paradigm	60
3.2	Research Design.....	61
3.3	Desktop Research (Secondary Data)	62
3.4	Selection of Sources.....	64
3.5	Data Analysis: Thematic Analysis.....	66
3.6	Methodological Framework: DECA (Describe, Evaluate, Consult, Act).....	68
3.6.1	Describe.....	68
3.6.2	Evaluate.....	68
3.6.3	Consult	69
3.6.4	Act	70
3.7	Conclusion	71
CHAPTER FOUR: THEORETICAL FRAMEWORKS – THE ETHICAL THEORY OF CONSEQUENTIALISM		72
4.0	Introduction.....	72
4.1	The ethical theory of Consequentialism	72
4.2	Definition and meaning of Consequentialism.....	74

4.3 Approaches of Consequentialism	75
4.3.1 Act Consequentialism	75
4.3.2 Rule Consequentialism.....	76
4.3.3 Motive Consequentialism.....	76
4.4 Different Types of Consequentialism	77
4.4.1 Utilitarianism.....	77
4.4.2 Hedonistic Consequentialism.....	78
4.4.3 Preference Consequentialism	78
4.4.4 Ethical Egoism (Individual Consequentialism).....	78
4.4.5 Negative Consequentialism.....	78
4.4.6 Ideal Consequentialism	79
4.5 Proponents of Consequentialism.....	79
4.6 Arguments for the Ethical Theory of Consequentialism	81
4.7 Arguments Against the Ethical Theory of Consequentialism.....	83
4.8 The Role of The Ethical Theory of Consequentialism in Guiding the Work	85
4.9 Conclusion	85
CHAPTER FIVE: ANALYSIS OF AI RECRUITMENT WITH THE ETHICAL THEORY OF CONSEQUENTIALISM.....	87
5.0 Introduction.....	87
5.1 Benefits of Using AI in Recruitment	88
5.1.1 Increased Efficiency in Recruitment	89
5.1.2 Reduced Human Bias	91
5.1.3 Cost Savings in Recruitment	92
5.2 Challenges of using AI in Recruitment.....	93
5.2.1 Algorithmic Bias and Discrimination in AI Recruitment	94
5.2.2 Lack of Transparency and Accountability in AI Recruitment	95
5.2.3 Dehumanization of the Recruitment Process	97
5.2.4 Over-Reliance on Technology	98
5.3 The Theory of Consequentialism Framework	99
5.3.1 Table 1: Ethical Assessment of Benefits of Using AI in Recruitment.....	99
5.3.2 Table 2: Ethical Assessment of Challenges of Using AI in Recruitment	101
5.8 Conclusion	102
CHAPTER SIX: SUMMARY, CONCLUSION, AND RECOMMENDATIONS.....	104
6.0 Introduction.....	104

6.1 Summary of the Study	105
Chapter One: Introduction.....	106
Chapter Two: Literature Review	106
Chapter Three: Research Method and Methodology	107
Chapter Four: Theoretical Framework: Consequentialism	108
Chapter Five: Analysis and Discussion.....	109
Chapter Six: Conclusion and Recommendations	110
6.2 Conclusion	111
6.3 Recommendations.....	112
6.3.1 Safeguard Human Dignity in Recruitment.....	113
6.3.2 Institutionalise Transparency and Human Oversight	114
6.3.3 Moderate Algorithmic Bias through Inclusive Data and Regular Audits	115
6.3.4 Build Ethical Capacity and Awareness	116
6.3.5 Reduce Over-Reliance through Governance and Accountability Frameworks	117
References	120

CHAPTER ONE: INTRODUCTION TO THE RESEARCH

1.0 Introduction

This study critically examines the ethical effects of using Artificial Intelligence (AI) in job recruitment and selection processes in South Africa through the lens of consequentialist ethics. AI has become a defining force in modern labour markets, transforming how employers identify, evaluate, and hire candidates. From automated résumé screening and psychometric analysis to predictive algorithms, AI-driven recruitment promises greater efficiency, objectivity, and consistency. Yet, despite these advantages, the ethical implications of automation in employment decision-making have sparked growing concern. Questions persist about bias, fairness, transparency, and accountability issues that have direct consequences for social justice and human dignity.

In the South African context, these concerns acquire moral urgency. The country's post-apartheid labour landscape remains deeply shaped by historical inequalities, high youth unemployment, and limited access to digital infrastructure. Consequently, any technology that influences employment outcomes must be ethically evaluated for its social consequences. A consequentialist approach provides a robust framework for this analysis, focusing on outcomes rather than intentions, that is, whether the use of AI in recruitment produces more benefit than harm for individuals, institutions, and society. This chapter introduces the research by outlining the background and motivation of the study, presenting the main and sub-research questions, defining the objectives, and summarising the methodological approach. It also provides a structural overview of the dissertation's organisation. The discussion situates the research within broader debates about ethics, technology, and governance in South Africa's evolving digital economy.

1.1 Background and Motivation

The integration of Artificial Intelligence (AI) into recruitment processes has become a global phenomenon, promising increased efficiency, cost reduction, and data-driven objectivity (Upadhyay and Khandelwal, 2020, p. 12). However, in South Africa, a nation grappling with structural unemployment (32.6% in Q3 2023) and persistent inequality rooted in apartheid-era policies (Statistics South Africa, 2023, p. 5), the adoption of AI recruitment tools presents unique ethical dilemmas. This study emerges against the backdrop of two conflicting imperatives: corporate demands for hiring automation to streamline processes, and legislative mandates for equitable employment under the Employment Equity Act (EEA) and Broad-

Based Black Economic Empowerment (B-BBEE) frameworks (South African Labour Guide, 2022, p. 8).

South Africa's labour market remains stratified along racial and socio-economic lines, with Black Africans disproportionately represented in low-wage sectors (Nkosi and Ruggunan, 2023, p. 18). AI systems trained on historical hiring data risk codifying these disparities, as evidenced by cases where algorithms penalised non-Western names or qualifications from historically Black universities (Bolukbasi et al., 2019, p. 9). This contradicts the transformative intent of EEA Section 15, which requires employers to "eliminate barriers to equitable representation" (Republic of South Africa, 1998, p. 12). The Protection of Personal Information Act (POPIA) further grants candidates the right to an explanation for automated decisions (Section 71), yet most AI recruitment platforms continue to function as "black boxes" (Binns, 2020, p. 3). A 2023 audit of four major South African corporations revealed that none could fully explain how their AI tools filtered applicants, citing vendor confidentiality (Datt et al., 2023, p. 14). This opacity undermines compliance with both POPIA and the Promotion of Administrative Justice Act (PAJA), which mandates procedurally fair decision-making.

Multinational AI recruitment tools often ignore South Africa's contextual realities, such as Recognition of Prior Learning (RPL). According to Mkhize (2023, p. 114), many skilled workers lack formal certifications due to apartheid-era educational policies, relying instead on informal or community-based experience. Algorithms often discount self-employment or volunteer work common in township economies (Dlamini, 2024, p. 24), while natural language processing (NLP) tools optimised for English underperform with African-language résumés (Abrahams et al., 2023, p. 8). These tensions culminate in the study's central research question: How does the use of AI in recruitment by South African companies affect ethical outcomes related to fairness, transparency, and legal compliance within the country's unique socio-legal context? The urgency of this inquiry is underscored by emerging regulatory shifts. In 2023, the Information Regulator issued draft AI Governance Guidelines emphasising "equitable design" (Mokgoro, 2023, p. 6). Meanwhile, the Commission for Employment Equity (CEE, 2023, p. 21) reported a 12% decline in racial transformation progress within AI-adopting sectors. By investigating these dynamics, this study contributes to the growing discourse on aligning global AI ethics frameworks with South Africa's redress-oriented labour policies.

The fourth industrial revolution has accelerated the integration of AI into various sectors, with human resources (HR) and recruitment being among the most significantly impacted.

Organisations worldwide are turning to algorithmic systems to streamline hiring processes, reduce administrative costs, and minimise human error. According to Chamorro-Premuzic et al. (2020, p. 78), AI recruitment tools can cut time-to-hire by up to 60% and enhance the quality of candidate selection through data-driven analysis. However, the use of these technologies has raised significant ethical questions regarding bias, accountability, and fairness. Globally, scholars such as Ajunwa (2020, pp. 12–13) and Bogen and Rieke (2020, p. 15) warn that AI systems can unintentionally reproduce discriminatory hiring patterns when trained on biased historical data. In South Africa, where employment equity and redress are constitutional imperatives, such risks have profound ethical implications. The Constitution of the Republic of South Africa (1996, section 9) guarantees equality and non-discrimination, while the Employment Equity Act (1998) aims to rectify past injustices in access to employment. The ethical challenge, therefore, is to ensure that AI recruitment aligns with these legal and moral commitments, advancing inclusion rather than exacerbating exclusion.

This study is motivated by the persistent tension between technological progress and moral responsibility. While AI offers potential benefits such as enhanced efficiency and fairness, its unintended consequences may undermine human dignity and social trust. As Driver (2015, pp. 97–98) explains, consequentialism evaluates moral rightness based on outcomes, the balance between good and harm generated by an action. This framework allows for an assessment of whether AI recruitment, as implemented in South Africa, genuinely enhances collective welfare or perpetuates harm through algorithmic bias, opacity, or exclusion. Moreover, South Africa's socio-economic context magnifies the ethical stakes. Persistent inequality, digital divides, and structural unemployment mean that access to fair recruitment processes is not merely a human resources concern but a moral and developmental imperative. A consequentialist analysis provides a lens through which to examine how AI affects various social groups, particularly those historically marginalised, and whether it contributes to or hinders the nation's pursuit of equitable development and social transformation. This study seeks to advance ethical governance by bridging the gap between AI innovation and human values. It aligns with global initiatives, such as the UNESCO Recommendation on the Ethics of Artificial Intelligence (2021), which advocate for technology that is transparent, accountable, and oriented toward human welfare. The motivation behind this study thus stems from both scholarly inquiry and civic responsibility to ensure that South Africa's digital transformation proceeds in a way that enhances, rather than erodes, collective well-being and justice.

1.2 Research Problem

The ethical challenges and consequences that arise when Artificial Intelligence (AI) is used in recruitment and selection processes in South Africa.

1.3 Key Research Question

What are the ethical implications of using AI in job recruitment and selection processes in South Africa?

1.4 Research Sub-Questions

1. What is the nature of traditional recruitment processes in South Africa?
2. How does AI-driven recruitment differ from traditional methods?
3. What are the ethical consequences of replacing traditional processes with AI in South Africa?

1.5 Key Research Objective

To evaluate the ethical implications of integrating AI into job recruitment and selection processes in South Africa.

1.6 Research Sub-Objectives

1. To analyse traditional recruitment processes in South Africa.
2. To investigate AI-driven recruitment mechanisms and challenges.
3. To assess the ethical consequences of AI adoption in recruitment.

1.7 Preview of Research Methodology and Method

This study adopts a qualitative, exploratory, and desktop-based research design to critically examine the ethical effects of Artificial Intelligence (AI) in job recruitment and selection within South Africa. Rather than collecting new empirical data, the research relies exclusively on secondary sources, including academic literature, policy documents, and legal frameworks. This approach, commonly referred to as desktop research, enables the synthesis and critical analysis of existing knowledge to illuminate the ethical tensions surrounding AI-driven recruitment (Travis, 2016, p. 1). A qualitative exploratory design is particularly suited to this inquiry, as it prioritises ethical interpretation, conceptual depth, and contextual understanding over statistical generalisability. As Lincoln and Guba (1985, p. 290) note, qualitative research seeks to explore meaning and complexity, making it ideal for under-researched and morally complex topics such as algorithmic hiring. This design allows the study to interrogate how ethical theories, especially consequentialism, apply to real-world technological practices.

Consequentialist ethics, which evaluates actions based on their outcomes rather than intentions, provides the normative framework for this research. The central aim is not to assess whether AI recruitment is technically effective, but whether its social and ethical consequences promote welfare, reduce harm, and uphold justice in South Africa's employment landscape.

The study draws on purposively selected secondary data from credible sources, including: Peer-reviewed journals in applied ethics, AI governance, and business ethics, South African legal instruments such as the Constitution (1996), Employment Equity Act (1998), Protection of Personal Information Act (2013), and Promotion of Administrative Justice Act (2000) and Global ethical frameworks such as UNESCO's Recommendation on the Ethics of Artificial Intelligence (2021) and the World Economic Forum's AI Governance Framework (2025). Only academically credible and policy-relevant materials were included. The purposive sampling strategy ensured that all sources directly addressed the intersection of AI, ethics, and employment, enabling a focused and coherent analysis of fairness, transparency, accountability, and human dignity. For data analysis, the study employed the DECA model Describe, Evaluate, Consult, and Act (Kretzschmar and Bentley, 2013, p. 5), a structured ethical decision-making framework aligned with consequentialist reasoning. This model guided the process of describing the ethical issues and their socio-legal context, evaluating the consequences of AI recruitment on moral and social welfare, consulting relevant philosophical and policy perspectives, and formulating recommendations to enhance ethical governance and mitigate harm. In parallel, thematic analysis (Braun and Clarke, 2019) was used to identify recurring ethical patterns across the literature, linking theoretical insights to practical implications. As the study involves no human participants, there were no concerns regarding consent or personal risk. Nonetheless, ethical integrity was upheld through rigorous citation, academic honesty, and adherence to the University of KwaZulu-Natal's Research Ethics Policy and international standards for responsible scholarship. In summary, this methodological framework integrates qualitative inquiry, desktop research, and normative ethical reasoning under a consequentialist paradigm. It provides a coherent structure for evaluating AI recruitment practices through the lens of welfare maximisation, harm reduction, and social justice, grounded in South Africa's unique socio-economic and regulatory context.

1.8 Preview of Theoretical Framework: Theory of Consequentialism

Consequentialism is an outcome-oriented ethical theory that evaluates the morality of actions based on their consequences. According to this framework, an action is morally right if it maximises positive outcomes such as welfare, happiness, or utility and wrong if it results in

net harm. This approach shifts ethical reasoning from intent to impact, urging decision-makers to weigh benefits against harms in pursuit of the greatest good. The most prominent form of consequentialism is utilitarianism, which seeks to “produce the greatest happiness for the greatest number.” Jeremy Bentham introduced the felicific calculus, a method for quantifying pleasure and pain by assessing factors such as intensity, duration, and scope. John Stuart Mill later refined this view, distinguishing between higher (intellectual and moral) and lower (bodily) pleasures. Mill famously asserted that “it is better to be a human being dissatisfied than a pig satisfied,” underscoring the moral superiority of higher-order goods. Modern philosophers have expanded consequentialism to address systemic and global ethical challenges. Peter Singer’s (2015) advocacy for effective altruism calls for evidence-based actions that maximise positive impact on humanity. Derek Parfit (1984) introduced intergenerational ethics, arguing that future generations have moral claims on present actions. Population ethicists such as Arrhenius, Bykvist, and Campbell have further explored how to weigh the welfare of different population sizes in policy decisions. These developments have led to justice-sensitive forms of consequentialism that consider not only aggregate utility but also fairness, equity, and institutional responsibility.

This theoretical lens is particularly relevant to technology ethics, where the focus shifts from developers’ intentions to the real-world effects of technological deployment. In the context of AI-driven recruitment, consequentialism demands a rigorous evaluation of both the benefits and harms experienced by all stakeholders. On the benefit side, AI recruitment tools promise: Efficiency gains: Rapid processing of large applicant pools, Improved matching: Predictive analytics to identify qualified candidates, Bias mitigation: Systems that ignore names or photos may reduce affinity bias, and Cost and time savings: Lower cost-per-hire and faster time-to-fill (Upadhyay and Khandelwal, 2020; Chamorro-Premuzic et al., 2020). These outcomes can enhance collective welfare by improving productivity, reducing unemployment, and offering more equitable access to opportunities (Driver, 2015). However, consequentialism also highlights significant ethical risks: Algorithmic bias: AI systems trained on biased historical data can replicate and amplify discrimination (Bolukbasi et al., 2019; Mehrabi et al., 2021), Opacity and accountability: Proprietary algorithms often lack transparency, making it difficult to challenge unfair decisions (Selbst and Barocas, 2018; Hunkenschroer and Luetge, 2022), and Dehumanisation: Reducing candidates to data points can erode dignity and alienate job seekers (Newlands, 2021; Binns, 2020). From a consequentialist perspective, these harms, especially when they result in systemic exclusion or psychological distress, must be weighed

against the promised benefits. If AI recruitment tools increase efficiency but deepen inequality or undermine trust, the net outcome may be ethically unjustifiable.

This ethical calculus is especially urgent in South Africa, where employment equity is both a constitutional imperative and a moral necessity. With unemployment at approximately 32.9% in early 2025 (Statistics South Africa, 2025), hiring practices directly affect social welfare. The Constitution (1996, s. 9), Employment Equity Act (1998), POPIA (2013), and PAJA (2000) collectively enshrine the rights to fairness, dignity, and just administrative action. Consequentialism aligns with these values by demanding that AI recruitment systems demonstrably promote inclusive employment and reduce structural disadvantage. In summary, consequentialism offers a structured moral framework for evaluating AI in recruitment. It compels stakeholders to: Measure outcomes empirically (for example, diversity of hires, candidate satisfaction, algorithmic fairness), implement safeguards (for example, bias audits, transparency protocols, human oversight), and Align AI practices with social justice and constitutional values. Only when AI recruitment demonstrably enhances collective welfare without reproducing historical injustices can it be ethically justified. Consequentialism thus bridges technological innovation and moral accountability, urging continuous reflection on whether AI serves the public good.

1.9 Significance of the study

This study makes a novel contribution to the interdisciplinary literature on AI ethics, recruitment, and moral philosophy. While prior work has documented how AI can improve hiring efficiency and potentially support diversity goals, scholars caution that algorithmic hiring also raises new ethical challenges, such as issues of fairness, transparency, and accountability. What is distinctive here is the application of a consequentialist ethical framework to AI-driven recruitment in the South African context. Consequentialist ethics judges actions by their outcomes, essentially that outcomes are all that matter, and people should act to produce the best overall consequences. By explicitly adopting this perspective, the dissertation adds a normative, outcome-oriented analysis to the literature. It extends discussions of algorithmic fairness, often couched in technical or deontological terms, by examining whether AI hiring practices maximize overall welfare and respect for human dignity. In doing so, the study bridges gaps between the growing empirical research on AI in HR efficiency, bias, and ethical theory. For example, recent meta-analyses also emphasize that transparency and fairness are now crucial and require human oversight. This study deepens that conversation by testing these claims against a moral theory of consequences. Overall, the study

enriches academic debates on AI ethics by integrating consequentialist moral reasoning into the specific domain of employment selection, a conjunction that is under-explored in both AI and business ethics literature.

The findings will have direct relevance for South African policymakers, HR professionals, and organizations that deploy AI in hiring. Employment equity and anti-discrimination laws are especially important. Under the Employment Equity Act (EEA 55 of 1998), employers must “eliminate unfair discrimination in any employment practice, including recruitment”. This means AI tools cannot indirectly disadvantage protected groups, such as race, gender, age, or disability. The study’s insights on how AI algorithms can produce unintended but harmful outcomes, as shown in cases like *Mobley v. Workday*, will help HR departments audit and adjust their tools. For instance, regular bias audits and human review mechanisms recommended by experts become practical steps to ensure AI compliance with both the letter and spirit of the EEA. Data privacy is another key concern. The Protection of Personal Information Act (POPIA, Act 4 of 2013) governs the processing of candidate data. POPIA “places strict duties” on employers to process personal information lawfully and securely. AI screening often involves collecting sensitive data, so HR teams must obtain clear consent, limit data use to stated purposes, and ensure secure storage. Critically, POPIA also recognizes a data subject’s right not to be subject to solely automated decisions without safeguards. As one commentary notes, affected candidates should have “sufficient information about how an AI decision was taken” and an opportunity to challenge it. The dissertation’s examination of AI decision-making can guide employers in creating transparent workflows that honour these legal requirements, for example, by explaining AI outputs to applicants or retaining human decision-makers at key stages.

For public-sector hiring, the Promotion of Administrative Justice Act (PAJA, Act 3 of 2000) applies. PAJA grants every person the right to administrative action that is “lawful, reasonable and procedurally fair” (Constitution sec.33). This implies that if a government body uses AI to select candidates, the process must be rational and non-arbitrary. In *Care phone (Pty) Ltd v Marcus NO*, the court held that any discretionary decision must have an “objective rational basis” linking the evidence to the outcome, otherwise it is arbitrary and unlawful. This study highlights that opaque or unvalidated AI systems risk failing this rationality test. Policymakers and regulators can utilize these insights to develop guidelines or regulations that mandate AI-based hiring systems to provide audit trails, clear criteria, and avenues for appeal, thereby integrating PAJA’s fairness requirements into technology governance. This research informs

the development of the emerging national AI strategy. South Africa's draft National Artificial Intelligence Policy Framework (2024) already identifies "transparency and explainability, fairness and mitigating bias, and human oversight" as three of its twelve key pillars. The study provides concrete content to these otherwise abstract pillars, for example, by suggesting how human oversight should be structured in an HR context or what bias mitigation looks like in practice. In summary, the study will provide evidence-based recommendations that help HR leaders and policymakers align AI hiring tools. With South Africa's legal and ethical standards, as well as the EEA, POPIA, PAJA, and the Constitution, and with international best-practice principles.

The study addresses deep societal concerns in post-apartheid South Africa. Employment practices are not value-neutral: fair hiring is a powerful tool for redressing historical inequalities and protecting human dignity. The Constitution explicitly recognizes the values of human dignity, equality, and non-discrimination as foundational to South African democracy. AI systems that encode historical bias could undermine these constitutional commitments by perpetuating the "apartheid-era" disparities in education and employment. Conversely, well-designed AI has the potential to help level the playing field by standardizing criteria and surfacing overlooked talent, for example, by matching candidates to B-BBEE and EE targets. This study's critical analysis will clarify how to make sure AI hiring strengthens, rather than subverts, our constitutional ideals of equality and dignity. Practically, the research contributes to fair hiring practices in a society with extremely high youth unemployment, which exceeds 40-45%, and skills gaps (Statistics South Africa, 2025). By proposing ethical guardrails such as bias audits and candidate redress mechanisms, the study promotes inclusive workforce participation and economic transformation. It reminds organisations that every candidate's dignity must be respected, even when using advanced technology. In doing so, it resonates with the African ethical principle of ubuntu, often summarized as "I am because we are," which implies that business practices should uphold community values, solidarity, and mutual respect. In summary, the study's significance lies in its potential to inform law and practice while advancing social justice. It fills a scholarly gap at the intersection of AI ethics and labour studies, providing South African employers and regulators with the conceptual and practical tools needed to use AI responsibly. By embedding respect for human dignity and equality into the algorithmic hiring process, the research aims to ensure that technological advancements in HR do not compromise fundamental human rights.

1.10 Definition of key concepts

1.10.1 Artificial Intelligence (AI)

Artificial Intelligence refers to the simulation of human intelligence processes by machines, especially computer systems. In the context of recruitment, AI systems utilise machine learning algorithms to automate and enhance various stages of the hiring process, including screening résumés, conducting initial interviews via chatbots, and predicting candidate success (CMS Law, 2024). These systems are trained on vast datasets, which is a critical point of ethical concern, as historical data often reflects and embeds societal biases (Pasipamire, 2024). Contextualized for critique to improve service delivery and efficiency; however, it introduces complex challenges regarding data privacy and the potential for algorithmic discrimination. As noted by Obasa (2023) and Maimela (2025), the responsible deployment of AI requires a comprehensive understanding of its technical capabilities and its socio-legal implications within the country's unique context.

1.10.2 Recruitment

Recruitment is the process of actively seeking out, finding, and hiring candidates for a job or position. It is a critical function of Human Resource Management (HRM) that aims to attract a qualified pool of applicants from which the best candidates can be selected (Ncube, 2025). The introduction of AI has fundamentally transformed this process, moving from traditional, human-centric methods to automated, data-driven systems. Cliff Dekker Hofmeyr (2025) states that, contextualised for the critique, the use of AI in recruitment in South Africa is rapidly increasing, but this automation must be carefully managed to ensure compliance with stringent anti-discrimination laws, such as the Employment Equity Act 55 of 1998. Also, Asoba and Mefi (2021, pp. 1-6) note that the ethical critique focuses on whether AI-driven recruitment systems genuinely broaden the talent pool or merely perpetuate existing inequalities by favouring candidates whose profiles align with historical, potentially biased, successful hires.

1.10.3 Ethical Recruitment

Ethical recruitment refers to the process of attracting, assessing, and selecting candidates in a manner that upholds fairness, transparency, and respect for human dignity. According to Breugh (2023, p. 45), ethical recruitment ensures that hiring decisions are based on merit and competence rather than on discriminatory factors such as race, gender, or socio-economic background. It involves adherence to moral and legal standards that promote justice and equality of opportunity throughout the recruitment cycle. In the context of Artificial

Intelligence (AI), ethical recruitment further demands the responsible use of algorithms to prevent bias, protect candidate privacy, and ensure accountability in automated decision-making (Ajunwa, 2020, pp. 12–13). Ethical recruitment thus integrates both procedural fairness, ensuring equal treatment, and substantive fairness, ensuring that outcomes do not reinforce social inequality. It aims to align technological innovation with human values, promoting trust between employers and job seekers while sustaining organizational legitimacy.

1.10.4 Algorithmic Bias

Algorithmic bias happens when an AI system generates outcomes that are consistently prejudiced toward or against certain individuals or groups. As Bolukbasi et al. (2019, p. 8) note, several varied factors can result in a biased model, from biased training data, poor model architecture, or the social assumptions that are encoded in the input features. In a recruiting setting, algorithmic bias can lead to candidates being rejected based on race, gender, language, educational attainment, and socio-economic status, all of which it seems plausible the authors would just as soon see purging from the human decision-making process, even when no explicit rule against it exists. This is particularly challenging in Stock and by extension in South Africa, where the past is penetrated by structural inequalities in qualifications, employment history, and institutional representation (Nkosi and Ruggunan, 2023, p. 18). In the absence of active bias detection and mitigation approaches, AI-powered recruitment systems pose a threat of enshrining existing inequalities under the banner of neutrality.

1.10.5 Fairness

Fairness in AI recruitment refers to the principle that an automated system should not produce outcomes that unjustly disadvantage or discriminate against individuals or groups based on protected characteristics such as race, gender, age, or disability (Cliff Dekker Hofmeyr, 2025). Asoba, Samson, and Mefi, Nteboheng (2025) highlights that achieving fairness is technically complex, as biases can be introduced at multiple stages, from the historical data used for training to the design of the algorithm itself. Contextualised for the critique, the concept of fairness is legally mandated in South Africa by the Employment Equity Act, which prohibits unfair discrimination. The critique must explore various definitions of algorithmic fairness (for example, demographic parity, equal opportunity) and assess whether current AI systems can meet the high standard of non-discrimination required to address South Africa's legacy of inequality and promote equitable employment practices.

1.10.6 Transparency and Explainability

Transparency and explainability are also pivotal requirements for such AI-based systems. Transparency is about how much employers, candidates, and regulators can learn about the workings and logic of algorithmically made recruitment decisions. Many such AI systems are ‘black boxes,’ resulting in minimal, if any, visibility into how or why an applicant has been shortlisted or rejected (Binns, 2020, p. 2). Explainability goes further than transparency by demanding that AI systems make understandable and meaningful the reasons for the outputs they produce (Wachter et al., 2018, 23). This, in terms of recruiting, means that not only should a candidate be able to get a rejection release, but also an explanation that explains which rejections impacted their score, and why. Explainability is also an essential ethical and legal demand in South Africa, as under the necessary frameworks of the Protection of Personal Information Act (POPIA) and the Promotion of Administrative Justice Act (PAJA). Absent this crucial information, people are prevented from challenging or even understanding decisions that have a dramatic impact on their ability to access employment, thereby compromising due process and confidence in the hiring process.

1.10.7 South Africa

South Africa provides a distinctive socio-economic and ethical landscape within which to examine the use of Artificial Intelligence (AI) in recruitment and selection. The country’s post-apartheid transformation agenda is rooted in constitutional values of equality, human dignity, and social justice, as articulated in the Constitution of the Republic of South Africa (1996). These principles guide all public and private institutions in addressing historical injustices and structural inequality. Consequently, any technological adoption in employment practices must align with the imperatives of fairness, inclusion, and non-discrimination. Economically, South Africa faces persistent challenges of high unemployment, particularly among youth and graduates, alongside widening inequality and skills mismatches (Statistics South Africa, 2024, p. 14). These dynamics heighten the ethical sensitivity surrounding AI-driven recruitment. While automation offers potential efficiency gains and improved matching between employers and job seekers, it also risks deepening exclusion if algorithmic systems replicate existing biases in education, language, or geography. Legally, South Africa has a robust framework governing labour and information ethics, including the Employment Equity Act 55 of 1998, the Labour Relations Act 66 of 1995, the Protection of Personal Information Act (POPIA), and the Promotion of Access to Information Act (PAIA). These statutes collectively emphasize equity, transparency, and accountability values that must guide AI deployment in recruitment.

As Mokgoro (2023, p. 9) argues, embedding ethical reasoning into technological innovation is essential to preserving public trust and ensuring that digital transformation supports, rather than undermines, South Africa's developmental goals. From an ethical perspective, South Africa's context demands a context-sensitive consequentialist approach, one that assesses not only efficiency outcomes but also broader societal impacts such as empowerment, fairness, and social cohesion. In this way, South Africa serves as both a testing ground and a moral benchmark for evaluating how AI can contribute to inclusive and equitable labour markets.

1.11 Research Structure

This dissertation is organised into six interrelated chapters, each contributing to a progressive ethical evaluation of Artificial Intelligence (AI) recruitment within the South African context. The structure is designed to move systematically from foundational concepts to applied ethical analysis and to practical recommendations. Each chapter builds upon the preceding one, ensuring theoretical, methodological, and analytical coherence. Collectively, the chapters aim to provide a comprehensive understanding of how consequentialist ethics can be used to assess the moral legitimacy of AI recruitment systems in a society grappling with inequality, unemployment, and technological transformation.

Chapter One: Introduction to the Research

The first chapter establishes the foundation of the study by introducing the research background, motivation, problem statement, aims, objectives, and significance. It discusses the growing adoption of AI in recruitment and selection processes and the associated moral ambiguities that arise from automation in decision-making. The chapter explains why evaluating AI recruitment through an ethical lens is essential, especially in South Africa, where structural inequality and employment equity remain pressing concerns. It also introduces the research methodology and outlines the justification for using consequentialism as the theoretical framework. By framing the study around outcomes that promote fairness, welfare, and justice, Chapter One lays the groundwork for an ethically informed examination of AI technologies in recruitment.

Chapter Two: Literature Review

Chapter Two offers a comprehensive review of global and South African literature related to AI, ethics, and employment. It explores existing academic debates on algorithmic bias, discrimination, transparency, accountability, and the dehumanisation of hiring processes. Drawing on key scholars such as Ajunwa (2020), Bogen and Rieke (2020), Chamorro-

Premuzic et al. (2020), and Zou and Schiebinger (2021), the review critically evaluates how AI influences fairness and inclusion in recruitment. It also highlights the limited South African-specific research in this field, demonstrating the need for a context-sensitive ethical analysis. The chapter concludes by identifying theoretical and empirical gaps that justify this study's focus and by establishing consequentialism as the appropriate ethical framework for evaluating the moral outcomes of AI recruitment.

Chapter Three: Research Method and Methodology

This chapter details the methodological design and approach employed in the study. It explains the choice of a qualitative, exploratory, desktop-based research design, which is well-suited for examining conceptual and normative ethical issues. Data were sourced from scholarly articles, legal documents, government reports, and policy frameworks such as the Employment Equity Act (EEA), the Protection of Personal Information Act (POPIA), and the Promotion of Administrative Justice Act (PAJA). The chapter justifies the use of purposive sampling to select relevant sources and introduces the DECA (Describe, Evaluate, Consult, Act) model as a systematic ethical analysis tool. Furthermore, it elaborates on methods used to ensure trustworthiness and credibility, following Lincoln and Guba's (1985) criteria of dependability, transferability, and confirmability. The methodological discussion establishes how consequentialist ethics is operationalised in this study by linking data analysis to moral evaluation of outcomes.

Chapter Four: Theoretical Framework – Consequentialism

Chapter Four explores consequentialism as the guiding ethical framework for this research. It defines consequentialism as an outcome-based theory asserting that the morality of an action is determined by its consequences rather than its intentions or adherence to rules (Driver, 2015). The chapter traces the historical evolution of consequentialist thought from classical utilitarianism, advanced by Bentham and Mill, to modern interpretations developed by de Lazari-Radek and Singer (2017), Mejía (2019), Segall (2021), and Arrhenius, Bykvist, and Campbell (2020). The discussion clarifies how consequentialism provides an effective ethical lens for evaluating AI recruitment because it prioritises measurable impacts such as welfare, fairness, and inclusion. The framework is contextualised within South Africa's moral and socio-political realities, where technological ethics must align with constitutional principles of equality, human dignity, and justice. This theoretical grounding prepares the reader for the analytical evaluation in the following chapter.

Chapter Five: Analysis and Discussion

This chapter applies the consequentialist framework to analyse the ethical implications of AI in recruitment. It focuses on five interrelated themes: operational efficiency, transparency and accountability, fairness and equity, candidate experience and human dignity, and broader societal consequences. Each theme is examined in relation to how AI-driven hiring systems affect welfare, justice, and moral legitimacy. The analysis integrates insights from scholars such as Chamorro-Premuzic et al. (2020), Hunkenschroer and Luetge (2022), and Segall (2021), while situating these within South Africa's unique employment and governance context. The chapter demonstrates that while AI can improve efficiency and reduce certain forms of bias, it also poses ethical risks such as algorithmic discrimination, opacity, and dehumanisation. Through a consequentialist lens, the chapter concludes that AI recruitment is ethically acceptable only when its use demonstrably increases collective welfare, safeguards fairness, and upholds human dignity.

Chapter Six: Conclusion and Recommendations

The concluding chapter synthesises the study's findings, reaffirming that the ethical acceptability of AI recruitment depends on the balance between its benefits and harms. It concludes that consequentialism provides a robust ethical framework for assessing technological innovations in recruitment by focusing on real-world outcomes rather than theoretical intentions. The chapter presents practical recommendations for policymakers, organisations, and developers to ensure ethical AI governance, data transparency, bias mitigation, and human oversight. It also identifies areas for future research, particularly concerning long-term and intergenerational effects of AI on employment equity and social cohesion. Ultimately, Chapter Six reinforces the central argument that technology must remain a means to human welfare, not an end, and that ethical governance grounded in consequentialist reasoning is essential to guide AI-driven recruitment towards justice, accountability, and collective well-being.

Overall, this structured approach ensures conceptual clarity and analytical depth. Each chapter builds logically upon the previous one, guiding the reader from theoretical foundations and methodological design to applied ethical analysis and practical recommendations. The organisation of the dissertation thus supports a coherent and comprehensive exploration of the ethical consequences of AI recruitment in South Africa, grounded in consequentialist ethics and oriented toward advancing fairness, human dignity, and social good.

1.12 Conclusion

This chapter has introduced the central focus of the study, a critical examination of the ethical effects of Artificial Intelligence (AI) in job recruitment and selection within the South African context. It has provided a comprehensive background and motivation for the research, highlighting how rapid technological advancement intersects with persistent structural inequalities, unemployment, and regulatory challenges rooted in South Africa's socio-historical realities. The chapter also articulated the main research question, sub-questions, and objectives, positioning the inquiry within the moral framework of consequentialist ethics, which evaluates actions by their outcomes rather than their intentions. Through this lens, the research seeks to determine whether AI-driven recruitment processes promote social welfare, fairness, and justice, or whether they risk deepening exclusion, bias, and opacity in employment practices. The qualitative, exploratory, and desktop-based methodological design was justified as appropriate for this philosophical and ethical analysis, as it integrates moral reasoning with empirical policy evaluation. The application of the DECA method (Describe, Evaluate, Consult, Act) and thematic analysis ensures that the study proceeds through a structured, reflective, and ethically grounded process of inquiry.

By outlining the research structure, this chapter has also demonstrated how each subsequent section builds toward answering the central question of whether AI recruitment is ethically defensible within South Africa's constitutional and socio-economic framework. The theoretical orientation established here, the consequentialist evaluation of outcomes, serves as the moral and analytical backbone of the dissertation. The next chapter, Chapter Two: Literature Review, develops this foundation further by critically engaging with existing academic and policy literature on AI, ethics, and recruitment. It examines global and South African scholarship to uncover prevailing debates on algorithmic bias, fairness, accountability, and the moral legitimacy of automation in employment. In doing so, the literature review identifies conceptual gaps that justify this study's focus on consequentialism as a lens for ethical evaluation. It also situates the research within the broader philosophical discourse on technology, governance, and justice, setting the stage for the theoretical and analytical developments that follow in subsequent chapters.

CHAPTER TWO: LITERATURE REVIEW

2.0 Introduction

The previous chapter provided an overview of the research problem and context, highlighting South Africa's history of employment discrimination and the rise of AI in the hiring process. It introduced the study's objectives and methodology, emphasizing the need to examine AI hiring tools through South Africa's legal and ethical framework. The chapter concluded by underscoring the importance of a comprehensive literature review to situate this inquiry, which this Chapter undertakes.

This review of the literature critically examines the ethical implications of Artificial Intelligence (AI) in recruitment and selection within the South African context. As AI-powered hiring tools become increasingly common around the world, they promise greater efficiency, objectivity, and scalability (Upadhyay and Khandelwal, 2020, p. 12; Chander, 2021, p. 45). However, their use also raises complex ethical and legal questions, especially in societies marked by structural inequality and historical injustice. The review addresses key aspects of this issue: defining AI in recruitment, distinguishing AI-driven from generic screening processes, understanding algorithmic bias, and emphasizing transparency and accountability. It also examines relevant legal frameworks, including the Employment Equity Act (EEA), the Protection of Personal Information Act (POPIA), and the Promotion of Administrative Justice Act (PAJA), as well as strategic approaches to AI (both short- and long-term, abstract and instrumental) and fundamental critiques of machine-led hiring. A critical concern is that most existing literature on AI in recruitment comes from the Global North and assumes more standardized labor markets and robust regulations. There is a gap in the ethical effects of using AI in job recruitment and selection in South Africa. This review is informed by leading scholars Sandra Wachter (2018) on algorithmic accountability and explainability, Cathy O'Neil (2018) on the dangers of automated "weapons of math destruction," and Reuben Binns (2020) on fairness and justice in machine learning, combined with South African labor studies to construct a critical lens for analyzing AI in recruitment.

2.1 Traditional Recruitment Process

The traditional recruitment process refers to the conventional methods organisations use to attract, screen, and select candidates before the introduction of digital and algorithmic technologies. This process typically involves human-led decision-making, relying on manual assessments, interviews, and paper-based applications. According to Armstrong and Taylor (2020, p. 214), traditional recruitment follows a linear sequence of stages: job analysis,

advertisement, shortlisting, interviewing, and final selection conducted primarily by human resource practitioners. The process aims to identify candidates whose skills, qualifications, and values align with the organisation's needs. Historically, the traditional recruitment process has been grounded in human judgment and interpersonal interaction. It allows recruiters to evaluate candidates' personalities, communication skills, and emotional intelligence through direct engagement (Cole, 2011, p. 85). This human-centred approach ensures a personalised assessment that considers not only technical competence but also cultural fit and potential for growth within the organisation. In the South African context, traditional recruitment methods have been widely used across both public and private sectors, particularly where human discretion and contextual understanding are essential (Nel et al., 2016, p. 112). Job advertisements under the traditional model are often disseminated through newspapers, internal postings, or career fairs, reflecting a localised and relationship-based hiring culture (Breugh, 2017, p. 67). Candidates typically submit physical applications or emailed documents, which are manually reviewed by HR departments. Interviews, whether structured or semi-structured, play a significant role in evaluating the suitability of applicants (Robertson and Smith, 2001, p. 449). The final hiring decision rests with managers or panels who exercise discretion informed by experience, intuition, and organisational norms.

The strength of traditional recruitment lies in its personalisation and transparency. Recruiters can provide immediate feedback, clarify ambiguities, and maintain two-way communication throughout the process. As Oaya, Ogbu, and Remilekun (2017, p. 32) note, the process allows for "contextual fairness" because recruiters can consider individual circumstances and qualifications that may not fit rigid templates. Moreover, human recruiters are capable of ethical deliberation, balancing efficiency with compassion, a skill that algorithmic systems struggle to replicate (Newlands, 2021, p. 92). However, Bogen and Rieke (2020, p. 15) argue that the human-led nature of traditional recruitment also introduces subjectivity and inconsistency. Decisions may be influenced by unconscious bias, nepotism, or stereotyping, undermining principles of fairness and equality. As Driver (2015, p. 98) argues, despite these limitations, traditional recruitment remains morally significant because it recognises the inherent dignity of candidates as autonomous individuals, a principle that aligns with both deontological and consequentialist ethics when the process produces outcomes that enhance collective welfare and respect for persons. In South Africa, the traditional approach remains particularly relevant due to the country's socio-economic and historical realities. Recruitment often occurs within communities where social networks and human recommendations play a

vital role in employment access, noted by Nkosi and Ruggunan (2023, p. 18). While digital transformation is advancing rapidly, many sectors, especially government, education, and small enterprises, continue to rely on traditional recruitment systems. As Statistics South Africa (2023, p. 5) states, these methods ensure inclusivity for candidates who may lack digital access or literacy, thereby preventing exclusion in a context still shaped by inequality and technological divides. Consequently, the traditional recruitment process, while not without flaws, has played a foundational role in maintaining fairness, accessibility, and human connection in employment practices. As AI and automation increasingly reshape recruitment systems, understanding the ethical implications of this shift requires acknowledging what may be lost, particularly human empathy, contextual reasoning, and the moral dimension of direct human engagement, qualities that traditional recruitment has long embodied.

2.1.1 The Nature of the Traditional Recruitment Process

The nature of the traditional recruitment process is fundamentally human-centric, characterised by personal interaction, professional judgment, and contextual evaluation. Unlike algorithmic or AI-driven recruitment, which relies on automated data processing, traditional recruitment depends on human discretion and relational assessment. According to Armstrong and Taylor (2020, p. 216), this approach rests on three interrelated pillars: personal engagement, experiential decision-making, and contextual fairness. Recruiters evaluate candidates not only on measurable qualifications but also on intangible attributes such as personality, motivation, and potential cultural fit within the organisation. This human-led nature provides depth and flexibility in assessment. As Cole (2011, p. 89) explains, the process allows recruiters to interpret non-verbal cues, emotional intelligence, and interpersonal skills dimensions that are critical in many job roles but difficult to quantify digitally. Moreover, Robertson and Smith (2001, p. 452) share that human recruiters can adapt their evaluation criteria based on evolving circumstances or organisational needs, ensuring that recruitment decisions remain responsive rather than rigidly rule-based.

In the South African context, the traditional recruitment process is further shaped by socio-economic realities, legislative imperatives, and cultural values. Recruitment often incorporates considerations of Employment Equity (EEA, 1998, p. 12) and Broad-Based Black Economic Empowerment (B-BBEE) compliance, ensuring that hiring contributes to social transformation. The process also reflects the spirit of ubuntu, a humanistic philosophy that values compassion, respect, and community interconnectedness by treating each candidate as

an individual worthy of dignity (Mkhize, 2023, p. 114). As Driver (2015, p. 97) notes, this moral dimension aligns closely with consequentialist ethics, which evaluates actions based on their ability to promote overall welfare and social harmony. Traditional recruitment methods typically involve structured steps such as job analysis, vacancy advertising, application collection, screening, interviews, and reference checks. While Breugh (2017, p. 71) mentions that each stage relies heavily on human interpretation and ethical judgment. Interviews are central to the process because they allow decision-makers to holistically evaluate a candidate's potential beyond the written CV. Additionally, traditional reference checks provide moral assurance regarding character and reliability, a feature often absent in AI-based assessments that focus narrowly on predictive metrics.

Another defining feature of the traditional process is transparency and accountability. Because human recruiters are solely responsible for their decisions, there exists a clear moral and procedural chain of accountability. As Nel et al. (2016, p.115) observe, recruiters can be questioned about their decisions, allowing candidates to seek clarification or lodge appeals, thereby upholding fairness. This personal accountability contrasts sharply with algorithmic recruitment, where decision-making may be opaque, and responsibility diffused across software systems and developers (Raghavan et al., 2020, p. 7). Despite the emergence of digital tools, Statistics South Africa (2023, p. 5) states that many South African organisations continue to rely on this traditional model because it accommodates contextual realities such as uneven digital literacy, limited internet access, and the need for face-to-face verification. Furthermore, Nkosi and Ruggunan (2023, p. 18) also mention that traditional recruitment reinforces relational trust between employers and candidates, a cornerstone of sustainable employment relationships in societies with histories of inequality and mistrust. In summary, the nature of traditional recruitment lies in its human-centeredness, ethical flexibility, and contextual responsiveness. It allows moral reasoning, empathy, and situational awareness to guide hiring decisions, qualities that consequentialist ethics recognises as central to producing morally positive outcomes. As AI recruitment technologies increasingly automate these processes, understanding the intrinsic moral and social value of human-led recruitment becomes essential for evaluating whether automation truly enhances or undermines welfare, fairness, and dignity in South Africa's employment landscape.

2.1.2 Stages of the Traditional Recruitment Process

The traditional recruitment process is a systematic, step-by-step method designed to attract, assess, and select the most suitable candidates for organisational roles. It emphasises human judgment, transparency, and procedural fairness principles that align closely with consequentialist ethics, which evaluate actions by their outcomes for all stakeholders. As Armstrong and Taylor (2020, p. 215) observe, each stage of the traditional recruitment process contributes to moral and practical outcomes such as equity, accountability, and human welfare, which are central to both ethical theory and South African labour governance.

2.1.2.1 Job Analysis and Planning

The process begins with a detailed job analysis, where the organisation identifies the specific skills, qualifications, and personal attributes required for a role. This stage provides the ethical foundation for fairness in recruitment, ensuring that hiring decisions are guided by objective job requirements rather than personal bias (Bowen, 2009, p. 30). A well-conducted job analysis also aligns with the Employment Equity Act (EEA 55 of 1998, p. 12) by promoting transparency and equal opportunity, thereby contributing to socially beneficial outcomes, an essential consequentialist aim.

2.1.2.2 Advertising the Vacancy

Once the job is defined, the vacancy is advertised through appropriate channels such as newspapers, career fairs, internal circulars, and community notice boards. Traditional advertising plays an important ethical role in promoting accessibility and inclusivity. As Nel et al. (2016, p. 120) note, wide and inclusive advertising maximises the pool of potential candidates and prevents exclusion of those without access to digital platforms. In South Africa, where digital divides persist, this inclusivity enhances fairness and social welfare by allowing marginalised communities to participate in employment opportunities.

2.1.2.3 Application and Shortlisting

During this stage, candidates submit applications, and human resource practitioners screen them based on predetermined criteria. Shortlisting involves balancing merit with the transformational goals of employment equity. Human evaluators can apply discretion in ways that promote redress and social justice, outcomes which automated systems may overlook (Nkosi and Ruggunan, 2023, p. 18). This discretion allows for the consideration of contextual

factors and the advancement of diversity, fairness, and empowerment, particularly important in a society working to overcome historical inequalities.

2.1.2.4 Interviewing and Assessment

The interview is the cornerstone of the traditional recruitment process, allowing for direct interaction between employer and candidate. It enables recruiters to assess intangible qualities such as communication, ethical judgment, motivation, and interpersonal skills, attributes that are vital to organisational harmony but difficult to quantify algorithmically (Cole, 2011, p. 92). Interviews also provide candidates with an opportunity to express themselves, fostering procedural fairness and respect for human dignity (Constitution of the Republic of South Africa, 1996, s. 10). The consequences of this engagement, mutual understanding, and transparency are ethically desirable as they enhance trust and satisfaction on both sides.

2.1.2.5 Reference and Background Checks

After interviews, reference checks are conducted to verify the accuracy of information provided by candidates and to confirm their professional credibility. According to Armstrong and Taylor (2020, p. 218), this stage ensures ethical accountability by preventing fraudulent representation and protecting the organisation's integrity. It also contributes to broader social welfare by ensuring that selected candidates are competent and trustworthy, capable of positively impacting both the organisation and society.

2.1.2.6 Selection and Job Offer

The final selection is made based on the cumulative evaluation of candidates' qualifications, interview performance, and background verification. Decision-makers apply human judgment that balances efficiency with fairness and contextual understanding. The chosen candidate receives a formal job offer outlining the terms of employment, rights, and responsibilities. This stage embodies procedural justice, as defined under the Promotion of Administrative Justice Act (PAJA 3 of 2000, s. 33), by ensuring that the final decision is rational, lawful, and transparent.

2.1.2.7 Onboarding and Induction

The recruitment process concludes with onboarding, during which new employees are introduced to the organisation's culture, policies, and ethical expectations. According to Saunders et al. (2018, p. 150), effective induction fosters long-term productivity, a sense of

belonging, and ethical alignment outcomes that contribute to collective welfare and institutional cohesion. Each stage of the traditional recruitment process carries moral significance, as it directly affects individual and societal well-being. When conducted with transparency, fairness, and inclusivity, these stages yield positive outcomes such as equitable employment, institutional trust, and social stability. Conversely, neglecting these principles through practices like nepotism, discrimination, or procedural opacity can result in ethically adverse consequences. As Driver (2015, p. 97) and Singer (2015, p. 40) suggest, the structured, human-led recruitment process exemplifies how procedural integrity supports ethically sound and socially beneficial employment outcomes.

2.1.3 Power of Traditional Recruitment Methods in South Africa

The power of traditional recruitment methods in South Africa lies in their ability to balance professional judgment with social understanding, fostering equitable and culturally responsive hiring practices. While digitalisation has reshaped global recruitment, traditional approaches remain deeply embedded in South Africa's employment landscape due to their adaptability, interpersonal character, and alignment with the country's moral and constitutional values. As Nel et al. (2016, p. 117) observe, many organisations, especially within the public sector and small-to-medium enterprises, continue to rely on human-led recruitment because it ensures inclusivity for individuals who may not have access to digital technologies or online platforms. One of the main strengths of traditional recruitment is its human element, which allows recruiters to assess candidates holistically. Cole (2011, p. 90) notes that through interviews, personal interactions, and community networks, employers can evaluate qualities such as emotional intelligence, ethical behaviour, and interpersonal competence attributes that artificial systems struggle to measure. This approach supports South Africa's ubuntu-based philosophy, which emphasises relational ethics and mutual respect (Mkhize, 2023, p. 114). In a consequentialist sense, such an approach generates morally beneficial outcomes by promoting welfare, dignity, and trust among candidates and recruiters alike (Driver, 2015, p. 97). Traditional recruitment methods also play a transformative role in advancing social justice. Since the Employment Equity Act (EEA, 1998, p. 12) mandates fair representation of historically disadvantaged groups, human recruiters can consciously apply affirmative action principles during candidate evaluation.

This active moral reasoning ensures that employment decisions consider both merit and redress, aligning with South Africa's constitutional vision of equality (Constitution of the

Republic of South Africa, 1996, section. 9). Automated systems, by contrast, often fail to interpret the socio-historical meaning of such policies, thereby risking a re-entrenchment of systemic bias Binns (2020, p. 56) states. Another source of power in traditional recruitment lies in transparency and trust. Human-led processes provide candidates with direct communication, feedback, and recourse, which enhances procedural fairness. As Armstrong and Taylor (2020, p. 218) note, traditional recruitment allows decision-makers to be held accountable for their choices, ensuring moral responsibility for both the process and the outcome. Also, Nkosi and Ruggunan (2023, p. 18) argue that in contexts where trust in institutions has been undermined by historical inequities and corruption, personal engagement in recruitment helps rebuild legitimacy and confidence in employment systems. Furthermore, traditional recruitment supports contextual and cultural awareness. Recruiters often understand the social and economic barriers facing applicants, such as limited access to formal education or work experience due to apartheid-era inequalities. This awareness enables ethically sensitive decision-making, where potential and lived experience are valued alongside formal qualifications (Dlamini, 2024, p. 24). Traditional recruitment methods can serve as a corrective mechanism against structural exclusion, ensuring that the hiring process contributes to social welfare rather than exacerbating inequality.

Despite their slower pace and susceptibility to human bias, traditional recruitment methods retain normative power because they reflect values of fairness, compassion, and accountability. These qualities embody the consequentialist ideal that ethical actions are those that maximize well-being and minimize harm across society observed by Driver (2015, pp. 97–98). Within South Africa’s socio-economic context, where employment opportunities are intrinsically linked to social justice, the traditional recruitment process remains not just an administrative mechanism but an ethical practice that safeguards dignity and equality. In summary, the enduring power of traditional recruitment in South Africa lies in its moral adaptability and its alignment with national ethical priorities. It integrates human empathy with institutional accountability, bridging the gap between efficiency and equity. As technological systems increasingly take over recruitment functions, understanding the moral significance of traditional methods becomes crucial to ensuring that innovation does not erode the foundational values of fairness, inclusivity, and respect for human dignity.

2.2 Limitations of the Traditional Recruitment Process

While traditional recruitment is valued for its human-centredness, transparency, and procedural fairness, it has significant limitations in complex labour markets such as South Africa’s.

Structural inequality, digital exclusion, and entrenched institutional bias mean that human-led processes do not always deliver equitable or efficient outcomes. As Binns (2020, p. 56) notes traditional methods can undermine access to opportunity, institutional trust, and organisational performance. Driver (2021, p.15) observes that the moral worth of any recruitment system depends on whether it advances welfare, fairness, and justice; in practice, traditional recruitment sometimes fails to meet these standards.

2.2.1 Subjectivity and Human Bias

One of the most widely acknowledged limitations of traditional recruitment is its susceptibility to human bias. Recruiters' personal preferences, stereotypes, or unconscious assumptions can influence decisions regarding which candidates are shortlisted or hired. Ajunwa (202, p.14) observes that such biases often manifest in the form of affinity bias, gender stereotyping, or racial prejudice factors that can unfairly disadvantage competent candidates from marginalised backgrounds. In South Africa, where historical inequalities and racial hierarchies still affect access to opportunity, these biases have particularly harmful social consequences. Such practices yield unjust outcomes by diminishing overall societal welfare and perpetuating exclusion rather than promoting fairness or equality.

2.2.2 Limited Efficiency and High Resource Costs

Traditional recruitment methods are often time-consuming, administratively burdensome, and financially costly. Manual application screening, in-person interviews, and reference checks often require significant human labour, delaying hiring timelines and increasing operational expenses. As Chamorro-Premuzic et al. (2020, p. 78) explain, inefficiency in recruitment not only affects employers' productivity but also generates broader economic costs by slowing down employment cycles. Inefficiency represents a moral shortfall because it reduces positive outcomes economic participation, productivity, and candidate access to opportunities, while increasing frustration and resource waste.

2.2.3 Lack of Standardisation and Inconsistent Decision-Making

Human recruiters often apply varying standards in evaluating candidates, influenced by personal judgment, organisational culture, or situational pressures. This inconsistency can result in arbitrary or subjective decision-making observed by Braun and Clarke (2019, p. 80). Without objective evaluation metrics, candidates with similar qualifications may receive different assessments, undermining procedural fairness. This lack of standardisation leads to

ethically undesirable consequences, unfairness, inefficiency, and reduced public trust in institutional hiring practices.

2.2.4 Limited Accessibility and Exclusionary Practices

According to Nkosi and Ruggunan (2023, p. 12), Traditional recruitment processes may inadvertently exclude qualified individuals from disadvantaged backgrounds who lack access to mainstream advertising platforms or professional networks. For example, rural candidates or those without digital literacy may struggle to apply for formal positions advertised primarily through online or urban channels. This exclusion perpetuates inequality and contradicts the ethical goal of maximising collective welfare. As the Constitution of the Republic of South Africa (1996, s. 9) notes in the South African context, such inequities undermine constitutional principles of equality and social justice.

2.2.5 Vulnerability to Nepotism and Corruption

Another major limitation of traditional recruitment is the risk of nepotism and favouritism. In some cases, recruitment decisions may be influenced by personal relationships, internal politics, or bribery rather than merit, noted by Nkosi (2023, p. 118). Such practices compromise institutional integrity and public confidence, producing outcomes that privilege a few while harming the majority, which seeks to maximise the overall good. Moreover, nepotism undermines employment equity legislation, including the Employment Equity Act 55 of 1998, which mandates fair and non-discriminatory employment practices.

2.2.6 Inefficiency in Data Management and Candidate Tracking

Traditional recruitment processes often lack the data-driven infrastructure required to manage, organise, and analyse large volumes of applicant information effectively. Without integrated digital systems, recruiters rely on manual record-keeping, fragmented spreadsheets, and paper-based documentation methods prone to human error, duplication, and oversight. As Tambe et al. (2019, p. 34) note, such inefficiencies can result in qualified candidates being overlooked, mismatched appointments, and delays in decision-making. In South Africa, where many organisations, particularly in the public sector and small enterprises, still operate with limited digital capacity, these inefficiencies are magnified. The absence of centralised applicant tracking systems (ATS) or automated screening tools means that recruitment teams may struggle to maintain accurate records, compare candidate profiles systematically, or follow up with applicants on time.

2.2.7 Psychological and Emotional Impact on Candidates

Traditional recruitment processes can impose significant psychological and emotional strain on candidates, particularly when communication is delayed, inconsistent, or absent altogether. Newlands (2021, p. 92) notes that the lack of feedback following interviews or application submissions often leaves candidates in a state of uncertainty, which can negatively affect their self-esteem, motivation, and trust in employers. This emotional toll is exacerbated in high-unemployment contexts like South Africa, where job opportunities are scarce, and competition is intense. Opaque recruitment practices such as unclear selection criteria, unexplained rejections, or prolonged silence can foster feelings of alienation, anxiety, and disempowerment among job seekers. These experiences are not merely administrative oversights; they carry ethical weight, as they affect the dignity and psychological well-being of individuals navigating an already unequal labour market.

2.3 Critique of the Traditional Recruitment Process

A critical evaluation of the traditional recruitment process reveals that, although it has long served as the foundation of organisational staffing, its structure and methods increasingly fall short of the ethical, social, and operational standards demanded by South Africa's evolving labour market. As Driver (2015, p. 97) observes, ethical systems must be judged by their real-world outcomes. When assessed through this lens, traditional recruitment can be ethically deficient, particularly when its consequences fail to maximise fairness, efficiency, and collective welfare. Despite its strengths, such as human judgment, contextual sensitivity, and procedural transparency, traditional recruitment often produces mixed moral outcomes. It is vulnerable to subjectivity, bias, nepotism, inefficiency, and inconsistency, all of which can undermine the goals of equitable employment and social justice. These shortcomings are especially pronounced in South Africa, where recruitment must serve not only economic imperatives but also transformative justice, redress, and constitutional commitments to equality. Moreover, traditional methods frequently lack the data management capacity and scalability required to handle large applicant pools, leading to administrative delays, mismatched appointments, and missed opportunities. The emotional toll on candidates, particularly when feedback is absent or processes are opaque, further diminishes the moral legitimacy of these systems. In contexts marked by high unemployment and historical exclusion, such inefficiencies and injustices carry profound ethical weight.

2.3.1 Ethical Limitations: Inequality and Exclusion

Traditional recruitment practices often reinforce structural inequalities, particularly in societies shaped by historical injustices such as South Africa. Despite their appearance of neutrality, these practices often reflect and perpetuate existing disparities in access to opportunities. As Nkosi and Ruggunan (2023, p. 15) observe, entrenched racial, socio-economic, and geographic inequalities continue to influence who is considered employable and how recruitment decisions are made. Processes that rely heavily on personal networks, informal referrals, or elite educational credentials tend to favour individuals from historically privileged backgrounds, while systematically marginalising those from under-resourced communities. For example, candidates from rural schools or historically Black institutions may be overlooked in favour of graduates from urban, formerly white universities, not because of merit, but due to implicit prestige hierarchies embedded in recruiter perceptions. Mill (1861, p. 22) highlights that these outcomes generate social harm by reducing access to employment for disadvantaged individuals, thereby failing the utilitarian standard of promoting the greatest good for the greatest number. In a society where employment is a key driver of dignity, mobility, and social inclusion, such exclusionary practices have far-reaching consequences not only for individuals but for broader societal cohesion and economic justice. Moreover, these practices often contravene the spirit and letter of South African labour law. The Employment Equity Act 55 of 1998 explicitly mandates that recruitment processes must eliminate unfair discrimination and promote equitable representation. Yet, in practice, traditional methods may inadvertently perpetuate exclusion and inequity, particularly when they fail to account for the socio-historical context of candidates' educational and professional trajectories.

2.3.2 Procedural Inefficiency and Lack of Scalability

Traditional recruitment methods are often constrained by procedural inefficiencies that delay hiring decisions, inflate operational costs, and limit scalability. Manual screening of applications, multiple rounds of in-person interviews, and subjective shortlisting processes consume substantial organisational time and resources (Chamorro-Premuzic et al., 2020, p. 78). These inefficiencies are particularly problematic in large-scale recruitment drives or sectors with high vacancy turnover, where speed and consistency are essential. In the South African context, where public sector departments and small to medium-sized enterprises frequently operate under resource constraints, such inefficiencies can have amplified consequences. Prolonged vacancies in critical roles such as healthcare, education, or municipal services not only reduce organisational productivity but also impede service delivery, affecting broader societal welfare. As Tambe et al. (2019, p. 34) note, delays in recruitment generate

opportunity costs, as qualified candidates may seek employment elsewhere, and organisations may miss strategic growth windows.

2.3.3. Lack of Transparency and Accountability

Transparency is a cornerstone of ethical recruitment, as it enables fairness, trust, and procedural justice. However, as Binns (2020, p. 56) argues, traditional recruitment processes often operate with minimal procedural clarity, leaving candidates uncertain about how decisions are made and why they were rejected. Feedback is rarely provided, and selection criteria are seldom disclosed, fostering perceptions of arbitrariness and exclusion. This opacity undermines accountability, as stakeholders, including applicants, oversight bodies, and the public, cannot verify whether recruitment decisions were rational, equitable, or legally compliant. Edwards and Veale (2018, p. 47) argue that opaque decision-making erodes procedural justice by preventing meaningful scrutiny and appeal. In such environments, even well-intentioned decisions may appear biased or corrupt, reducing institutional legitimacy and diminishing social welfare.

2.3.4 Emotional and Psychological Impact

Traditional recruitment processes can impose significant emotional and psychological burdens on candidates, particularly when communication is delayed, unclear, or absent. Newlands (2021, p. 92) argues that prolonged silence after interviews, vague rejection notices, and a lack of constructive feedback contribute to stress, self-doubt, and disillusionment among applicants. These effects are not merely incidental; they reflect deeper ethical concerns about how candidates are treated within institutional systems that wield considerable power over their livelihoods and self-worth. In South Africa's high-unemployment environment, where job scarcity intensifies competition and heightens the stakes of each application, these emotional costs are amplified. For many candidates, especially first-generation graduates or those from historically disadvantaged communities, the recruitment process is not only a professional hurdle but a deeply personal and existential one. When that process is marked by opacity, silence, or perceived unfairness, it can erode confidence, discourage future applications, and reinforce feelings of exclusion.

2.3.5 Lack of Data-Driven Insight and Adaptability

Traditional recruitment methods rely on subjective human judgment, often in the absence of systematic data analytics or evidence-based evaluation tools. As Dlamini and Patel (2023, p.

7) observe, many organisations continue to depend on intuition, informal criteria, and anecdotal experience when assessing candidates and reviewing recruitment outcomes. This reliance on non-empirical methods limits the ability to track diversity metrics, detect bias patterns, or evaluate procedural fairness across recruitment cycles. Without a robust data infrastructure, organisations struggle to answer critical ethical and operational questions: Are historically disadvantaged groups being fairly represented? Are certain departments or roles consistently reproducing exclusion? Are recruitment outcomes aligned with the organisation's stated values and transformation goals? In the absence of such insight, harmful or ineffective practices may persist unexamined, undermining both institutional credibility and moral progress.

2.3.6 Nepotism and Corruption Risks

Traditional recruitment methods, particularly those that rely on personal networks, informal referrals, or opaque internal processes, are especially vulnerable to nepotism and corruption. These practices compromise the principle of meritocracy, where positions should be awarded based on competence and fairness, and instead introduce subjective or politically motivated decision-making. As Nkosi (2023, p. 118) notes, such dynamics undermine social justice by privileging those with access to influence while excluding equally or more qualified candidates without such connections. When recruitment decisions are shaped by who one knows rather than what one knows, the consequences are ethically and operationally damaging. They produce unequal outcomes, erode public trust in institutions, and violate the constitutional and legislative commitments to fair and non-discriminatory employment practices. This is particularly concerning in South Africa, where recruitment is expected to contribute to transformative justice and redress historical patterns of exclusion. Bogen and Rieke (2020, p. 15) argue that favouritism in hiring not only harms excluded candidates but also weakens organisational integrity, reduces productivity, and fosters a culture of entitlement rather than accountability. Corruption in recruitment can lead to the appointment of underqualified individuals, misallocation of public resources, and long-term reputational damage.

2.3.7 Ethical Interpretation through Consequentialism

Viewed through the lens of consequentialism, the traditional recruitment process falls short of meeting the moral criterion of producing the greatest overall good. While it upholds values such as human judgment, interpersonal engagement, and contextual sensitivity, its real-world outcomes are often biased, inefficient, and unequal, resulting in tangible social and economic harm. These shortcomings are especially concerning in South Africa, where recruitment carries

both a constitutional mandate and a moral responsibility to advance transformation, fairness, and equality (Constitution of the Republic of South Africa, 1996, s. 9). Consequentialist ethics, as articulated by Arrhenius, Bykvist, and Campbell (2020, p. 118), require that ethical evaluation consider both present and future impacts on collective welfare. Traditional recruitment, by maintaining structural inequities and failing to adapt to evolving social needs, does not fulfil this obligation. Its continued reliance on subjective judgment, opaque procedures, and exclusionary practices undermines the goal of maximising well-being across society. In this light, the critique of traditional recruitment establishes both a moral and practical rationale for exploring AI-assisted recruitment systems. If ethically designed and responsibly implemented, such systems have the potential to enhance efficiency, fairness, and transparency, while reducing bias and improving access to opportunity. However, this transition must be guided by ethical safeguards, ensuring that technological innovation does not replicate or deepen existing injustices.

2.4 Artificial Intelligence (AI) in Recruitment

The rise of Artificial Intelligence (AI) has fundamentally transformed how organisations attract, assess, and appoint talent. Within the context of the Fourth Industrial Revolution, AI-driven recruitment systems are increasingly deployed to automate tasks traditionally performed by human resource (HR) professionals, including candidate sourcing, résumé screening, interview scheduling, and performance prediction (Upadhyay and Khandelwal, 2020, p. 12). The appeal of AI in recruitment lies in its promise to increase efficiency, reduce costs, and enhance objectivity through data-driven insights. According to Chamorro-Premuzic, Winsborough, Sherman, and Hogan (2020, p. 78), AI-based recruitment tools can shorten hiring timelines by up to 60 percent while improving the accuracy and consistency of candidate evaluations. These technologies function by analysing large datasets to identify patterns associated with successful hires. Using natural language processing (NLP), machine learning (ML), and predictive analytics, AI systems evaluate applicants' qualifications, experiences, and competencies as observed by Bogen and Rieke (2020, p. 15). Algorithms trained on historical hiring data can automatically shortlist candidates, assess cultural fit, and even conduct pre-recorded video interviews using facial and voice recognition analytics (Raghavan et al., 2020, p. 6). This shift marks a paradigm transformation in human resources, moving recruitment from intuition-based decision-making to evidence-based, algorithmically supported processes.

In South Africa, the adoption of AI in recruitment is shaped by dual imperatives: economic efficiency and compliance with labour transformation policies, including the Employment Equity Act (EEA) and the Broad-Based Black Economic Empowerment (B-BBEE) framework. Organisations increasingly view AI as a tool for mitigating human bias, promoting fairness, and aligning hiring decisions with equity goals, noted by Nkosi and Ruggunan (2023, p. 24). However, despite its potential benefits, AI recruitment raises profound ethical and legal challenges, particularly in a society marked by historical inequality, digital divides, and high unemployment. Scholars caution that algorithms trained on biased historical data may reinforce existing social hierarchies, thereby contradicting South Africa's constitutional and legislative commitments to equality and non-discrimination (Binns, 2020, p. 3; Dlamini, 2024, p. 24). The integration of AI into recruitment thus represents both an opportunity and a moral test. On one hand, it offers the possibility of more transparent, consistent, and scalable hiring practices; on the other, it risks perpetuating systemic exclusion if ethical safeguards are not rigorously implemented.

As Driver (2015, pp. 97–98) argues, the moral legitimacy of any system depends not on its intentions but on its consequences. Applying this consequentialist perspective, AI recruitment must be evaluated according to whether it genuinely enhances social welfare by improving fairness, accessibility, and efficiency or whether it creates new forms of digital discrimination and opacity that undermine public trust. In the South African context, AI recruitment cannot be understood merely as a technological innovation. It must be assessed as an ethical and social intervention with far-reaching implications for justice, equity, and human dignity. The challenge lies not only in adopting advanced tools but in ensuring that these tools are aligned with constitutional values and transformative goals, thereby contributing to a more inclusive and accountable employment landscape.

2.4.1 Overview of AI in Recruitment

Artificial Intelligence in recruitment refers to the application of machine learning algorithms, predictive analytics, and natural language processing to automate and improve various stages of the hiring process. These technologies aim to assist HR professionals in decision-making by providing objective, data-driven insights that enhance both the speed and accuracy of candidate selection, noted by Upadhyay and Khandelwal (2020, p. 13). The recruitment process typically involves multiple stages: sourcing, screening, interviewing, and selection, and AI has applications in each. AI-powered platforms can scan online databases, job boards, and social

networks such as LinkedIn to identify potential candidates that match job descriptions, thereby expanding recruitment reach observed by Chamorro-Premuzic et al., (2020, p. 78). At the screening stage, AI systems evaluate resumes and cover letters using natural language processing to detect relevant keywords and competencies. Predictive models are then employed to estimate a candidate's likelihood of success in a particular role based on historical performance data (Raghavan et al., 2020, p. 8). For interviews, AI can conduct asynchronous video assessments that analyse facial expressions, tone, and speech patterns to infer personality traits or emotional intelligence, though these methods remain ethically controversial due to privacy and cultural sensitivity concerns noted by Newlands (2021, p. 92). After selection, AI tools assist in onboarding and performance prediction by identifying training needs and career development opportunities.

Globally, organisations such as Unilever and IBM have pioneered AI recruitment systems, reporting improved diversity outcomes and shorter hiring cycles noted by Bogen and Rieke, (2020, p. 15). However, these technologies are not without risk. AI systems often inherit the biases of their training datasets, leading to discriminatory outcomes if left unchecked. In one notable case, Amazon discontinued an AI recruitment tool after discovering that it systematically downgraded resumes containing words associated with women, such as "women's chess club," as observed by Zou and Schiebinger (2021, p. 324). Such examples illustrate that while AI promises objectivity, its ethical soundness depends entirely on the data quality and governance frameworks guiding its use. Within South Africa, the adoption of AI recruitment tools remains uneven, limited to large corporations and multinational subsidiaries with access to advanced HR technologies observed by Datt et al (2023, p. 14). Yet, there is growing interest among local organisations, particularly as they seek to align digital transformation strategies with the National AI Policy Framework (2024). The challenge, however, is ensuring that imported AI systems, often trained on Western data, are adapted to South Africa's socio-linguistic and cultural context. For instance, natural language processing tools designed for English may underperform when evaluating resumes written in African languages or non-standard dialects noted by Abrahams et al (2023, p. 8). The ethical evaluation of AI in recruitment depends on its outcomes, whether it produces fairer, more transparent, and inclusive hiring systems. The use of AI should serve collective welfare rather than merely organisational efficiency. If designed responsibly, AI can promote fairness by standardising decision-making, improving accessibility for candidates in remote areas, and reducing unconscious bias. Conversely, if deployed carelessly, it can automate exclusion, deepen

inequalities, and erode human dignity. Therefore, while AI recruitment represents a crucial step toward modernising human resource management, its integration must be guided by ethical principles, continuous auditing, and alignment with South Africa's broader goals of equity, transformation, and social justice.

2.4.2 The Nature of Artificial Intelligence in Recruitment

Artificial Intelligence (AI) in recruitment refers to the capacity of computer systems to perform cognitive functions traditionally associated with human decision-making, such as learning, reasoning, and problem-solving. As Upadhyay and Khandelwal (2020, p. 13) explain, within the hiring process, AI operates through algorithms capable of processing and interpreting large volumes of applicant data, recognizing patterns, and generating predictive insights that inform selection decisions. Bogen and Rieke (2020, p. 15) further argue that these systems utilize technologies such as machine learning (ML), natural language processing (NLP), and predictive analytics to automate functions that once required human judgment, including résumé screening, interview analysis, and candidate scoring. At its core, the nature of AI in recruitment is probabilistic rather than deterministic. Rather than making fixed judgments, AI systems rely on statistical correlations derived from training data to predict outcomes such as the likelihood of a candidate's success in a role (Raghavan et al., 2020, p. 7). As Zou and Schiebinger (2021, p. 324) observe, these predictions are extremely sensitive to the quality and diversity of data inputs. If training data reflect historical bias or exclusion, the system may replicate these inequities even when designed to be objective. Thus, AI recruitment systems do not inherently possess moral or ethical reasoning; instead, they mirror the moral logic or lack thereof embedded in the data and human decisions used to construct them.

consequentialism offers a fitting theoretical lens for understanding the nature of AI in recruitment. Driver (2015, pp. 97–98) evaluates the morality of actions based on their outcomes, focusing on whether they enhance welfare and reduce harm. Applying this framework to AI recruitment highlights the need to assess whether algorithmic decisions improve fairness, accessibility, and efficiency, or whether they exacerbate discrimination and exclusion. AI's moral neutrality, its lack of intrinsic intent, means that its ethical legitimacy must be judged solely by its real-world consequences within the social and institutional context in which it operates. In South Africa, this moral context is particularly significant. Mejía (2019, pp. 312–314), in his theory of institution-sensitive consequentialism, argues that ethical assessment must account for the institutional structures and social systems that shape outcomes.

Given South Africa's entrenched inequality and historical labour segregation, the use of AI in recruitment can have amplified consequences for marginalised groups. If AI systems replicate structural disadvantage through biased data or exclusionary design, they may perpetuate injustice even without malicious intent.

Yet, as Chamorro-Premuzic et al. (2020, p.78) note, when developed and deployed responsibly, AI recruitment can support equitable access to employment by reducing subjective human bias, promoting transparency, and improving the match between candidates and job opportunities. Therefore, the nature of AI in recruitment is best understood as dual-faceted: technically innovative yet ethically contingent. Its potential to transform hiring practices depends not only on computational accuracy but also on its alignment with moral and societal values. The ethical question thus becomes: Do AI recruitment systems produce outcomes that advance social welfare, fairness, and dignity? The answer depends on continuous ethical evaluation, institutional oversight, and data governance practices that ensure AI serves humanity rather than undermines it.

2.4.3 AI-Driven Recruitment Systems

AI-driven recruitment systems represent the practical application of Artificial Intelligence within the hiring process, using automated technologies to enhance or replace traditional human decision-making in candidate assessment. Tambe, Cappelli and Yakubovich (2019, p. 34) argue that these systems incorporate a range of digital tools and algorithms designed to streamline recruitment, from job advertising and candidate sourcing to interviewing and final selection. At the technical level, AI recruitment systems operate through several interconnected components: data collection, data analysis, and decision automation. In the data collection stage, AI systems gather information from resumes, online profiles, and digital assessments. This data is then analysed using machine learning algorithms to identify patterns of success linked to previous employees or organisational performance indicators. As Bogen and Rieke (2020, p. 15) mentioned that the final stage, decision automation, involves ranking, scoring, or shortlisting candidates based on predictive models of job suitability. Some systems go further by incorporating predictive personality analytics or AI-driven video interviews, which use facial expression and voice tone analysis to evaluate emotional intelligence or cultural fit (Newlands, 2021, p. 92). The strength of AI-driven recruitment lies in its efficiency, scalability, and capacity to reduce administrative workload. Chamorro-Premuzic et al. (2020, p. 78) note that AI tools can process thousands of applications within minutes, lowering costs and time-

to-hire. Furthermore, when properly calibrated, these systems can improve objectivity by reducing human bias arising from fatigue, prejudice, or subjective interpretation (Binns, 2020, p. 3). However, their growing autonomy introduces complex ethical and legal challenges. Raghavan et al. (2020, p. 8) argue that the lack of transparency in algorithmic processes can make it difficult to determine how decisions are made or to whom accountability belongs when errors or discriminatory outcomes occur.

In South Africa, these concerns are heightened by socio-economic disparities and legal frameworks that mandate fair and equitable hiring. The Employment Equity Act (1998) prohibits any employment practice that results in unfair discrimination, while the Protection of Personal Information Act (POPIA, 2013) requires that automated decision-making respect candidates' right to transparency and explanation, as the Republic of South Africa (1998, p. 12) and POPIA (2013, s. 71) state. Therefore, organisations deploying AI recruitment systems must ensure that algorithms are auditable, interpretable, and consistent with constitutional values of dignity, equality, and freedom. From a consequentialist standpoint, the ethical legitimacy of AI-driven recruitment systems depends on their aggregate outcomes for both individuals and society. Positive consequences include operational efficiency, expanded access to job opportunities, and potentially fairer candidate evaluations. Negative consequences include algorithmic bias, loss of human empathy in decision-making, and erosion of procedural fairness. As Segall (2021, p. 133) argues, justice-oriented consequentialism requires moral scrutiny of structural outcomes, even when harm is unintended. In practice, this means that South African institutions must evaluate not only whether AI hiring tools increase productivity but also whether they promote justice and inclusivity. AI-driven recruitment systems embody both promise and peril. They represent technological progress capable of transforming employment landscapes, yet their ethical impact hinges on how they are designed, monitored, and governed.

2.4.4 Definition and Scope of Artificial Intelligence in Recruitment

Artificial Intelligence (AI) in recruitment refers to the simulation of human cognitive functions such as learning, reasoning, and problem-solving through computational systems that assist or replace elements of the hiring process (Upadhyay and Khandelwal, 2020, p. 13). Within human resource management, AI is applied to analyse large volumes of candidate data, identify patterns, and generate predictive insights about job suitability and performance. Tambe, Cappelli, and Yakubovich (2019, p. 34) note that AI recruitment technologies encompass a wide spectrum of applications, including automated résumé screening, skill matching,

candidate engagement chatbots, and predictive analytics for workforce planning. The scope of AI in recruitment is both technical and ethical. Technically, it spans data-driven systems that collect, process, and interpret structured and unstructured data to support decision-making (Binns, 2020, p. 3). Ethically, it involves evaluating how such automation affects fairness, transparency, and human dignity in employment practices. Chamorro-Premuzic et al. (2020, p. 78) identify three broad domains of AI recruitment: Cognitive automation, which replaces repetitive administrative tasks; Predictive analytics, which forecasts candidate success, retention, or fit; and Interactional AI, which enhances communication between employers and applicants.

In the South African context, the definition and scope of AI in recruitment must be understood against a backdrop of socio-economic inequality and legal obligations for fairness. The Employment Equity Act 55 of 1998 prohibits discrimination in hiring, while the Protection of Personal Information Act 4 of 2013 (POPIA) regulates the use of candidate data. As noted in the Republic of South Africa (1998, p. 12), these frameworks enshrine constitutional imperatives of equality and dignity, which AI systems must uphold. The ethical and operational scope of AI recruitment in South Africa is broader than in many developed economies. It must align not only with goals of efficiency and innovation, but also with transformative justice and constitutional compliance. This expanded scope demands that the evaluation of AI recruitment extend beyond technical functionality to its broader social outcomes. As Driver (2015, pp. 97–98) argues, consequentialism holds that moral worth depends on the goodness of outcomes. Therefore, defining AI in recruitment necessarily includes its potential consequences for social welfare, whether it expands access, reduces bias, and enhances collective well-being, or conversely, reinforces exclusion and inequality. The scope of AI recruitment is as much moral as it is technological. It encompasses all tools, processes, and outcomes that influence the ethical quality of hiring decisions within the South African socio-legal landscape. Understanding this dual nature is essential for ensuring that AI systems contribute to fair, inclusive, and dignified employment practices.

2.4.5 Applications of Artificial Intelligence in Recruitment

The application of Artificial Intelligence (AI) in recruitment encompasses a wide array of technologies that automate, optimise, or augment decision-making across the hiring lifecycle. These applications typically span five key stages: sourcing, screening, assessment,

interviewing, and selection, each offering distinct operational advantages and ethical considerations.

2.4.5.1 Sourcing

In the sourcing phase, AI-powered tools such as programmatic advertising platforms and social media crawlers identify and target potential candidates by analysing online profiles and behavioural data. Upadhyay and Khandelwal (2020, p. 13) argue that these systems can increase the diversity and volume of applicant pools while reducing recruiter workload. However, Zou and Schiebinger (2021, p. 324) caution that sourcing algorithms may perpetuate digital bias if they rely on datasets that overrepresent certain demographics, thereby narrowing access for underrepresented groups.

2.4.5.2 Screening and Assessment

During screening and assessment, AI systems perform résumé parsing, keyword analysis, and predictive matching to shortlist candidates who meet predefined job criteria. Natural Language Processing (NLP) models assess applicants' skills and experiences by comparing them with historical hiring data (Bogen and Rieke, 2020, p. 15). While these tools improve efficiency and reduce manual bias, Bolukbasi et al. (2019, p. 8) warn that they may also introduce algorithmic discrimination when training data reflect historical inequities. For example, if past hiring practices favoured male engineers, the AI system might inadvertently downgrade female applicants with similar qualifications.

2.4.5.3 Interviewing and Psychometric Testing

AI is increasingly applied in video interviewing and psychometric testing, where machine-learning algorithms analyse facial expressions, speech tone, and body language to assess personality traits or cultural fit. Newlands (2021, p. 92) highlights the risk of dehumanising candidates by reducing interpersonal evaluation to data points. However, Tambe, Cappelli, and Yakubovich (2019, p. 34) argue that such systems can enhance objectivity when designed with diverse datasets and human oversight.

2.4.5.4 Selection and Decision-Making

At the final stage, predictive analytics and cognitive computing tools assist recruiters in evaluating long-term candidate potential and job retention likelihood. These systems forecast which applicants are most likely to thrive in specific roles or organisational environments. As

Singer (2015, p. 40) notes, the moral legitimacy of these applications depends on whether they maximise positive outcomes for both individuals and organisations, namely, fair hiring, productivity gains, and social inclusion.

2.4.5.5 South African Context

In South Africa, the practical application of AI recruitment technologies is expanding rapidly. Firms such as Giraffe and Graylink (2024) use algorithmic matching to address youth unemployment and streamline hiring. These examples illustrate AI's dual role as both an economic enabler and a social instrument: it can democratise job access while simultaneously demanding vigilant governance to prevent discriminatory effects. Each AI application must be assessed not only for operational efficiency but also for its ethical consequences, specifically, how it affects candidate dignity, fairness, and equality. As Segall (2021, p. 133) argues, moral evaluation must account for systemic outcomes, not merely individual acts. Thus, the ultimate measure of success for AI applications in recruitment is whether they contribute to the collective welfare of society and align with South Africa's transformative constitutional code.

2.5 Legal and Regulatory Compliance in Recruitment

The adoption of Artificial Intelligence (AI) in human resource management is advancing at a pace that often outstrips the development of corresponding legal frameworks both globally and within South Africa. This rapid integration raises significant legal and ethical concerns related to fairness, discrimination, privacy, and accountability in recruitment practices. In South Africa, the Employment Equity Act 55 of 1998 (EEA) prohibits both direct and indirect discrimination in recruitment and mandates affirmative action for historically disadvantaged groups. As Nkosi (2023, p. 112) argues, if AI hiring systems systematically exclude marginalised candidates, whether through biased training data, insufficient contextual calibration, or indirect exclusion mechanisms, they risk breaching this legislative mandate. Such exclusions not only violate statutory obligations but also undermine the country's broader goals of redress and transformation. Similarly, the Protection of Personal Information Act 4 of 2013 (POPIA) imposes strict conditions on how personal data is collected, processed, and used, particularly in the context of automated decision-making. Roos (2021, p. 115) notes that POPIA grants individuals the right to be informed when automated systems are employed and the right to object to such processing. However, these rights are often undermined when AI systems operate opaquely or are shielded by proprietary secrecy, making it difficult for candidates to understand or challenge algorithmic decisions.

The Promotion of Administrative Justice Act 3 of 2000 (PAJA) further reinforces the requirement for fair and just administrative action, mandating that affected individuals be provided with adequate reasons for decisions that impact their rights, including access to employment. In the context of AI recruitment, this implies that candidates must be able to contest and comprehend the basis of algorithmic outcomes, a standard that many current systems fail to meet. Internationally, the European Union's General Data Protection Regulation (GDPR) includes a right to explanation concerning automated decisions. Yet, as Edwards and Veale (2018, p.46) argue, enforcement of this provision remains inconsistent, with organisations often providing minimal human oversight or invoking proprietary algorithmic secrecy to avoid accountability. These gaps reveal that legal compliance alone is insufficient to ensure ethical recruitment. In the South African context, where constitutional values of social justice, equality, and dignity are paramount, AI recruitment systems must be designed for contestability, transparency, and ethical alignment. Legal frameworks provide a foundation, but ethical AI recruitment requires proactive governance that ensures fair treatment, data protection, and human oversight. The integration of AI into recruitment must be guided not only by regulatory compliance but by a commitment to ethical responsibility. Systems must be evaluated for their real-world impact on candidate rights, institutional accountability, and societal welfare. In this regard, consequentialist ethics offers a valuable lens emphasising that the moral legitimacy of AI recruitment depends on whether it produces outcomes that enhance fairness, protect dignity, and promote inclusive opportunity.

2.5.1 South African Legislation on Employment and AI Recruitment

South African employment law is deeply rooted in the constitutional principles of social redress, equality, and human dignity, forming the ethical foundation upon which recruitment practices must be evaluated. The integration of Artificial Intelligence (AI) into recruitment processes directly engages three key legislative instruments: the Employment Equity Act (EEA), the Protection of Personal Information Act (POPIA), and the Promotion of Administrative Justice Act (PAJA). Each of these statutes is transformative, impacted by the application of AI in hiring, particularly in relation to fairness, transparency, and accountability. The Employment Equity Act 55 of 1998 seeks to eradicate unfair discrimination and promote equal opportunities and affirmative action in the workplace. Section 15 of the Act mandates proactive initiatives to transform employment demographics, ensuring equitable representation of designated groups, including Black South Africans, women, and persons with disabilities. As Nkosi (2023, p. 112) argues, any AI system that disproportionately discriminates against

these groups, whether by design or inadvertent bias, may be said to violate this mandate. Importantly, legal review must consider nuanced forms of discrimination. Historically, neither courts nor society may have explicitly recognised that a function or process treats certain groups differently; however, they would observe that the outcomes appear skewed in ways that raise concern. For example, if an AI application consistently filters out résumés from non-Western educational institutions or township-based work histories, it may perpetuate structural exclusion, even if the algorithm appears neutral on the surface.

The Protection of Personal Information Act 4 of 2013 (POPIA) further regulates how personal data is collected, processed, and used, especially in automated decision-making contexts. POPIA grants individuals the right to be informed when AI systems are used and the right to object to such processing. These provisions are critical in ensuring that AI recruitment systems operate with transparency and consent, particularly when decisions affect access to employment. The Promotion of Administrative Justice Act 3 of 2000 (PAJA) reinforces the requirement for fair and just administrative action, mandating that individuals be provided with adequate reasons for decisions that impact their rights. In the context of AI recruitment, this implies that candidates must be able to understand and contest algorithmic outcomes, a standard that many proprietary systems currently fail to meet. South African legislation demands that AI recruitment systems be designed and implemented in ways that uphold constitutional values. This includes preventing direct and indirect discrimination; Ensuring data protection and informed consent; and providing procedural fairness and accountability in decision-making. Failure to meet these standards risks not only legal non-compliance but also ethical failure, particularly in a society committed to transformative justice and inclusive development. AI recruitment must therefore be governed by both technical precision and moral responsibility, ensuring that innovation serves the broader goals of equality, dignity, and social welfare.

2.5.1.1 The Protection of Personal Information Act (POPIA) (Act No. 4 of 2013)

The Protection of Personal Information Act (POPIA) (Act No. 4 of 2013) is South Africa's primary legislative framework governing the lawful processing of personal information. The Act establishes comprehensive rules to ensure that data processing respects individuals' privacy, autonomy, and dignity principles closely aligned with ethical governance and fairness in recruitment. POPIA mandates that personal data must be collected and processed lawfully, transparently, and for explicit purposes related to the legitimate functions of the organisation. As Roos (2021, p. 115) notes, the Act includes specific provisions addressing transparency in

automated processing, data access, and rectification rights, and significantly, the right of individuals not to be subject to fully automated decision-making without appropriate safeguards or human oversight. In the context of AI-driven recruitment, compliance with POPIA is critical. Many AI recruitment systems automatically capture and process sensitive categories of data, including biometric identifiers, for example, facial recognition inputs, psychometric assessments, and inferred behavioural or personality traits derived from curriculum vitae or digital interactions. Such data collection, when conducted without explicit consent or adequate disclosure, risks violating the principles of data minimisation, purpose limitation, and lawful processing enshrined in the Act. From a consequentialist ethical standpoint, non-compliance with POPIA has far-reaching moral and social implications. The harmful outcomes, such as privacy violations, discrimination, or exclusion based on algorithmic profiling, undermine collective welfare and erode public trust in technology. Conversely, ethical adherence to POPIA fosters transparency, fairness, and respect for individual dignity, thereby producing positive societal consequences that align with the broader aims of justice and inclusion. Therefore, organisations deploying AI in recruitment must not only meet the legal requirements of POPIA but also uphold its ethical spirit, ensuring that automated decision-making enhances human welfare rather than diminishing it. Incorporating human oversight, transparent data practices, and informed consent mechanisms represents both a legal obligation and a moral duty under consequentialist ethics.

2.5.1.2 The Employment Equity Act (EEA) (Act No. 55 of 1998)

The Employment Equity Act (EEA) (Act No. 55 of 1998) is a cornerstone of South Africa's legislative framework for promoting equality, fairness, and inclusivity in employment. The Act aims to eliminate unfair discrimination and ensure equitable representation across all occupational levels. It obliges employers to implement affirmative action measures to redress historical imbalances and create workplaces that reflect the country's demographic diversity. According to Dupper (2019, p. 142), the EEA enshrines both a legal and moral commitment to social justice, requiring that employment practices, whether human or automated, advance substantive equality rather than perpetuate systemic disadvantage. In the context of AI recruitment and selection, the EEA assumes renewed significance. Algorithmic systems, if not professionally designed and monitored, can inadvertently reproduce historical patterns of exclusion embedded in training data or decision-making models. As Ajunwa (2020, p. 1673) warns, the use of biased datasets in automated recruitment risks contravening the EEA's objectives by perpetuating structural inequalities under the guise of objectivity. To comply with

the Act, employers must therefore ensure that their AI systems undergo continuous auditing, validation, and human review to detect and mitigate discriminatory outcomes. From a consequentialist perspective, the ethical legitimacy of AI recruitment under the EEA depends on its outcomes, whether it enhances fairness, inclusivity, and welfare across diverse social groups. If AI systems demonstrably expand access to opportunities and support equitable hiring practices, their deployment contributes to positive moral and societal consequences. Conversely, if automation entrenches bias or exclusion, the resulting harm outweighs its utilitarian benefits, rendering its use ethically unjustifiable. The EEA demands that technological innovation align with the moral imperatives of equality and justice. Ethical governance requires that AI-driven recruitment not only comply with the legal provisions of the Act but also embody its transformative vision to create workplaces that reflect dignity, fairness, and shared prosperity for all South Africans.

2.5.1.3 The Labour Relations Act (LRA) (Act No. 66 of 1995)

The Labour Relations Act (LRA) (Act No. 66 of 1995) governs employment relationships in South Africa, promoting fair labour practices and protecting employees from unjust treatment. It provides a legal framework for collective bargaining, dispute resolution, and procedural fairness in employment decisions. According to Grogan (2022, p. 56), the LRA's central ethical principle lies in upholding fairness and dignity in all employment-related processes, ensuring that workers are not treated as mere instruments of production but as participants with inherent moral worth. In the era of AI-driven recruitment, the principles of the LRA face new challenges. Automated systems that screen, shortlist, or reject candidates without adequate human oversight may infringe on the procedural fairness required by the Act. As Keller and Aplin (2024, p. 9) observe, algorithmic decision-making often lacks transparency, making it difficult for applicants to contest or appeal unfair outcomes. This opacity can lead to violations of employees' and candidates' rights to fair treatment and due process. Furthermore, the absence of human reasoning in AI-mediated decisions risks eroding the trust and dialogue that the LRA seeks to promote between employers and employees. From a consequentialist perspective, ethical compliance with the LRA requires evaluating the outcomes of AI recruitment practices in terms of their effects on justice, trust, and welfare. As Singer (2015, p. 40) argues, moral responsibility involves acting in ways that maximise well-being and minimise harm. If AI tools contribute to efficiency while preserving fairness, transparency, and employee dignity, their use aligns with consequentialist ethics and the spirit of the LRA. However, if automation diminishes workers' rights, excludes potential employees, or

undermines procedural justice, then its moral and legal legitimacy is compromised. The LRA underscores that technology must serve human values, not replace them. Ethical recruitment through AI requires continuous human oversight, grievance mechanisms, and transparent communication channels that ensure accountability. By integrating these principles, organisations can harmonise technological innovation with South Africa's broader constitutional commitment to fairness, equity, and respect for human dignity.

2.5.1.4 The Promotion of Administrative Justice Act (PAJA) (Act 3 of 2000)

The Promotion of Administrative Justice Act (PAJA) (Act No. 3 of 2000) enshrines the constitutional right to lawful, reasonable, and procedurally fair administrative action. It ensures that individuals are treated justly by any authority exercising public power or performing administrative functions. According to Hoexter (2021, p. 45), PAJA gives practical effect to Section 33 of the Constitution, guaranteeing that all administrative decisions, including those involving automated or algorithmic systems, are transparent, rational, and subject to human accountability. In the context of AI recruitment, PAJA becomes particularly relevant when technology is used in the public sector or by private bodies performing functions that affect citizens' rights, such as hiring for state-owned enterprises. The Act mandates that affected persons must be given adequate notice, an opportunity to make representations, and apparent reasons for decisions. When AI systems make or influence hiring choices without providing explainable reasoning, so-called "black-box" algorithms, they undermine these procedural rights. As Binns (2020, p. 3) notes, opaque decision-making processes challenge principles of fairness and legality by preventing candidates from understanding why they were rejected or how their data influenced the outcome. From a consequentialist ethical perspective, the moral legitimacy of AI recruitment under PAJA depends on its social and human outcomes rather than its technological efficiency. As Driver (2015, pp. 97–98) and Singer (2015, p. 40) argue, actions are morally justified only when they maximise welfare and minimise harm. If AI recruitment enhances fairness, consistency, and access to employment opportunities, its use supports collective welfare. However, if it leads to unjust exclusion, lack of recourse, or diminished trust in public institutions, it generates ethically negative consequences. In South Africa, where fairness, equity, and dignity are central to constitutional governance, non-transparent AI hiring processes may also erode public confidence in administrative justice. For instance, if an AI system used by a municipality or government department rejects candidates without explanation, it may violate both Section 33 of the Constitution and Section 5 of PAJA, which require reasoned and reviewable administrative decisions. Therefore, employers must ensure

that AI-driven systems are auditable, interpretable, and capable of providing meaningful explanations for decisions. Consequentialism thus supports the integration of human oversight and procedural safeguards in AI-assisted hiring. Transparent design, regular audits, and the inclusion of appeal mechanisms not only promote legal compliance but also produce better moral outcomes, ensuring that technological innovation strengthens, rather than weakens, democratic accountability and human dignity in recruitment.

2.5.2 Comparative World Ecstasy Monitor

The World Ecstasy Monitor was not introduced to provide a systematic comparison of NFI data and DPs, but there is some merit in briefly comparing some of the results across the datasets. Legislation around the world is still trying to deal with the question of what role goes to algorithmic decision-making (Mendis, 2024). The European Union’s General Data Protection Regulation (GDPR) is frequently highlighted as the model marker. GDPR introduces data protection principles, such as data minimisation and purpose limitation, and grants individuals the right to explanation in automated decisions. Under GDPR Article 22, individuals have the right not to be subject to a decision based solely on automated processing, including profiling, which produces legal effects concerning them or similarly significantly affects them. In theory, this opens the door to challenging AI-based hiring decisions. However, as Edwards and Veale (2018, p. 46) point out, implementation of this requirement is often inconsistent and perfunctory. Organisations may skirt these obligations by using algorithms with minimal human oversight or by invoking proprietary protections. Under GDPR, individuals are not entitled to a detailed explanation of how specific AI decisions are made. Rather, as Wachter et al. (2018, p. 25) note, the regulation requires only general descriptions of the logic used, an approach many consider insufficient for meaningful accountability. This leaves candidates with limited recourse, even in jurisdictions with robust data protection frameworks.

2.5.2.1 Implication for South Africa

The enforcement challenges associated with the European Union’s General Data Protection Regulation (GDPR) serve as a cautionary tale for South Africa, where regulatory capacity remains limited, and AI recruitment tools are often imported “off the shelf.” These systems frequently arrive with embedded global biases and design assumptions that are misaligned with South Africa’s legal, cultural, and socio-economic realities. Most international platforms do not customise their systems to comply with local legislation, such as the Protection of Personal Information Act (POPIA), or labour policies like the Employment Equity Act (EEA). This creates an authority vacuum, in which companies can assert technological neutrality while

reproducing inequality through opaque and unaccountable algorithmic processes. As Selbst et al. (2019, p. 56) argue, AI governance must be understood as a structural issue, not merely a matter of technical compliance or regulatory checklists. In this view, ethical AI recruitment requires more than adherence to legal norms; it demands systemic safeguards that address the institutional and historical dynamics shaping hiring outcomes. These include Ethics by design principles that embed fairness and inclusion into system architecture; Algorithmic audits that detect and correct discriminatory patterns; and Data justice frameworks that ensure representativity, transparency, and accountability in data use. In the South African context, where employment is a key driver of redress and transformation, the failure to localise and ethically govern AI recruitment tools risks entrenching structural exclusion under the guise of innovation. Imported systems that are not adapted to South Africa's constitutional values, including equality, dignity, and social justice, may inadvertently undermine the very goals that employment law seeks to advance. Therefore, the implication for South Africa is clear: legal protections such as POPIA, EEA, and PAJA must be complemented by proactive governance mechanisms that ensure AI systems are not only technically functional but also ethically defensible. This includes local oversight, contextual calibration, and public accountability, ensuring that AI recruitment contributes to inclusive development rather than replicating global inequities.

2.6 Gaps in Literature

Despite an expanding body of research on AI in the context of recruitment worldwide, existing studies are strongly biased towards high-income countries with relatively homogeneous labour markets, strong institutional governance, and robust digital infrastructure (Raghavan et al., 2020, p. 10; Bogen and Rieke, 2018, p. 5). These studies assume widespread education, internet access, and formal employment structures that do not characterize much of the Global South. South Africa offers a quite different context: stark socio-economic inequality, racialized opportunities, multiple official languages, and a large informal economy. A significant gap exists in the incomplete understanding of how AI systems process applications written in African languages or hybrid dialects. Linguistic exclusion, in part, results from the constraints of most NLP tools (Abrahams et al., 2023, p. 8). That is one of the areas that tend to be overlooked, whether the AI recruitment technology that you use does or does not comply with South Africa's legal obligations in terms of the Employment Equity Act, POPIA, and PAJA. Little research addresses the question of whether it may be safe to deploy an AI system that could indirectly still discriminate against vulnerable groups. Truly little is known about how

algorithmic systems designed to scan curricula vitae and process formal assessments systemically sift out candidates from townships and the informal workforce, who often lack formal paper qualifications but possess valuable experiential knowledge (Dlamini and Patel, 2023, p. 9). However, these functional and legal concerns are not the primary focus of this study. The central gap addressed by this study concerns the ethical implications of AI-driven recruitment and selection in South Africa, examined from a consequentialist perspective. Specifically, there has been limited inquiry into how the real-world consequences of AI hiring technologies impact historically marginalized groups in post-apartheid South Africa, where employment equity is not only a legal obligation but also a moral imperative. Existing global frameworks typically emphasize deontological ethics, abstract fairness metrics, or technical mitigation strategies without examining the practical outcomes of AI hiring tools in contextually diverse and historically unequal societies.

2.7 Conclusion

The current research has taken a close look at the increasing adoption of Artificial Intelligence (AI) in recruiting and its ethical, legal, and social considerations. As much value as AI can provide in terms of operations, decreased time to hire, improved data analytics, and similar practices., this value comes with significant risk. Topping the list are algorithmic bias, opacity, and regulatory noncompliance, followed by the devaluation of human and cultural context in candidate assessments. Through global academia, scholars have posed warnings about AI systems replicating structural inequality, especially when those technologies become trained on biased historical data or designed in sociopolitical environments unlike the ones in which they are utilized. Most of this literature originates from high-income countries of the Global North and is based on homogeneous populations and well-functioning digitalized structures. Consequently, not enough attention has been paid to how AI works (and fails) in developing, post-colonial, and very unequal labour markets, such as the one South Africa has.

The South African experience is that employment equity is not a matter of moral or ethical option, but of constitutional and legislated right of action. Such tools that inadvertently discriminate against race, gender, language, education, or socio-economic status can be in breach of key laws such as the Employment Equity Act, POPIA (the Protection of Personal Information Act), and the PAJA (the Promotion of Administrative Justice Act). And beyond legality, the thoughtless use of AI could dilute the country's post-apartheid aspirations of social redress and transformation. The dominant algorithmic systems are often inconsiderate of the unique labour characteristics of South Africa; multilingualism, participation in the informal

sector, non-linear career trajectories, and educational disparities maintain systemic injustice. These are not challenges on the margins; they are organizational unpleasant facts. Failing to account for them in the design and application of AI recruitment systems facilitates exclusion by design, making worse the very disparities we typically think these tools will help to redress.

In this context, the present research is timely and relevant. By framing research within consequentialist ethical theory, we can consider AI in recruitment based on more than limited technical performance or theoretical declarations of fairness, but the actual effects, notably, on populations who have been historically excluded from the formal economy. Many of these movements to produce more contextually embedded AI ethics have been missing from global AI ethics discourse, and the focus on local laws and minority cultural values, as well as lived experiences of South African job seekers, has the potential to provide a critique that is less global in its scope and much more grounded. Finally, this study takes debates about fairness out of the abstract and examines the tangible effects of automated hiring on a society that continues to struggle with inequality. It hopes to more connect global conversations about AI ethics with South Africa's redress-based labour environment and has a theoretical and practical commitment to the future of ethical hiring. By synthesizing theoretical perspectives and empirical findings, this chapter has established a solid foundation for the research, identifying gaps in current knowledge, particularly in the application of consequentialist reasoning to AI recruitment ethics in South Africa. These insights inform the design and focus of the subsequent research, guiding the methodological choices needed to systematically investigate the ethical outcomes of AI in recruitment. The next chapter, Chapter 3, will outline the research method and methodology employed in this study. It will detail the qualitative exploratory design, the purposive selection of secondary sources, and the analytical framework used to evaluate the ethical consequences of AI in recruitment, thereby providing a clear roadmap for addressing the research questions identified in this study.

CHAPTER THREE: METHOD AND METHODOLOGY

3.0 Introduction

AI-driven recruitment and selection systems present complex ethical challenges, including issues of fairness, transparency, and privacy in hiring decisions. In the South African context, these concerns are heightened by the absence of specific AI regulations, raising risks of bias or discrimination if such tools are not carefully governed. Examining the ethical implications of AI-based recruitment thus requires a research approach that can capture nuanced, context-dependent insights. The preceding Chapter 2 examined the theoretical and empirical foundations of these ethical challenges. It explored global and South African perspectives on algorithmic bias, transparency, and accountability in AI recruitment systems. The chapter also situated these debates within ethical frameworks, particularly consequentialism, and reviewed how key concepts such as fairness, justice, and inclusion are applied in AI-driven hiring. In doing so, Chapter 2 established the conceptual basis for understanding how automated decision-making can both advance and undermine equity in employment, especially in a post-apartheid socio-economic landscape. Building on that foundation, this chapter outlines the methodology employed to investigate these issues, detailing the study's philosophical paradigm, research design, data collection and analysis methods, and the application of the DECA (Describe, Evaluate, Consult, Act) ethical decision-making framework. It also discusses the measures taken to ensure research quality, trustworthiness, and ethical integrity. Qualitative inquiry was chosen because it is well-suited to exploring complex social phenomena and the meanings ascribed to them. As Creswell and Poth (2018) observe, qualitative research can be envisioned as an "intricate fabric" of many threads, held together by general assumptions and interpretive frameworks. In keeping with this vision, the methodology for this study was carefully woven to support a deep, contextual examination of AI recruitment ethics in South Africa.

The sections that follow describe how an interpretivist paradigm, an exploratory qualitative design, desktop research for data collection, purposive sampling, and thematic analysis were used to address the research questions. The use of the DECA framework is explained as a guiding tool for ethical analysis, and issues of trustworthiness, ethical research practice, and study limitations are critically discussed. This reflective and rigorously structured methodology provides the foundation for analysing how AI-based hiring tools intersect with ethical principles and the South African socio-legal context. Following this overview, it is essential to articulate the philosophical foundation that underpins this study. Every methodological choice

reflects certain assumptions about how reality can be known and understood. Because the ethical challenges of AI recruitment are embedded in human interpretation, institutional context, and social meaning, this research aligns with an interpretivist paradigm. The interpretivist lens enables a nuanced understanding of how individuals, organizations, and systems construct, and experience ethical issues related to algorithmic hiring. It emphasizes meaning making over measurement, privileging depth, and context rather than generalization. The next section, therefore, outlines the philosophical paradigm guiding this study, explaining how interpretivism supports a qualitative, exploratory approach to examining the ethical implications of AI in recruitment and selection within South Africa.

3.1 Philosophical Paradigm

This study is grounded in an interpretivist constructivist paradigm, which posits that reality is socially constructed and that understanding emerges from interpreting the meanings that people or texts attribute to phenomena. An interpretivist paradigm seeks to understand the subjective world of human experience, acknowledging the existence of multiple realities shaped by context (Creswell and Poth, 2018, p. 24; Lincoln, Lynham, and Guba, 2018, p. 214). In the realm of ethical implications of AI recruitment, an interpretivist approach means recognizing that concepts such as fairness or accountability may be perceived differently by various stakeholders and within diverse cultural or organizational contexts. Rather than seeking singular, objective truths, the research embraces a relativist ontology of multiple truths and a subjectivist epistemology of knowledge co-created by the researcher and the data (Mertens, 2020, p. 21).

Under this paradigm, the researcher's role is to interpret and make sense of qualitative data such as texts, policies, and prior studies to uncover patterns of meaning related to AI ethics in hiring. This approach is appropriate because the study deals with normative and context-dependent questions, for example, how and why certain AI-driven hiring practices might be seen as ethically problematic in South Africa. Such questions benefit from an exploratory, understanding-oriented lens rather than a purely positivist one (Bryman, 2016, p. 384). The interpretivist stance also aligns with the need to consider the social, cultural, and legal context: South Africa's history and values around equality and diversity are likely to shape what is deemed "ethical" in AI recruitment (Chilisa, 2020, p. 38). By adopting an interpretive paradigm, the research can sensitively account for these contextual factors.

Creswell and Poth (2018, p. 25) note that qualitative researchers typically operate within interpretive frameworks that hold the study together conceptually. For this inquiry, the interpretivist paradigm provides that conceptual backbone. It guided the formulation of research questions focused on understanding perceptions and implications, the choice of methods, qualitative and flexible, and the analytical strategy, thematic interpretation rather than statistical analysis. It also implies a reflexive approach: the researcher continually reflected on how her own perspectives and values might shape the interpretation of findings, an important consideration since ethical analysis is inherently value-laden (Tracy, 2020, p. 3). In summary, the philosophical orientation of this study is interpretive, prioritizing depth of understanding and contextual meaning-making as the foundation for examining the ethical dimensions of AI in recruitment within the South African context.

3.2 Research Design

Given the exploratory nature of the inquiry, the study employed a qualitative, exploratory research design. The goal was to obtain rich, detailed insights into the ethical dimensions of AI-based recruitment rather than to test a hypothesis or measure prevalence. A qualitative design is appropriate because the research questions centred on ethical implications, perceptions, and frameworks require nuanced analysis of language, concepts, and context that quantitative methods could not adequately capture (Creswell and Poth, 2018, p. 44; Braun and Clarke, 2022, p. 5). The design can be described as descriptive and interpretive, aiming to describe the phenomena, such as types of ethical issues arising from AI recruitment, and interpret their significance within the South African setting. Concretely, the research was carried out as a desk-based study using secondary data sources. This means no new data were gathered from human participants; instead, the analysis drew upon existing literature, documents, and reports relevant to AI in recruitment and ethics.

The design is thus akin to a qualitative content analysis or documentary analysis of available information (Bowen, 2009, p. 27). This approach was chosen for several reasons. First, the topic is new and interdisciplinary, so pertinent information is dispersed across academic papers, industry white papers, legal texts, and policy documents. A desk research design allowed the integration of these multiple sources to form a comprehensive view (Flick, 2018, p. 134). Second, given the ethical focus, many relevant insights come from previous analyses, cases, or theoretical discussions. Mining these sources can illuminate potential issues and principles without needing to experiment directly on an AI system or conduct a survey (Merriam and Tisdell, 2016, p. 19). Third, practical considerations such as time, access, and ethical feasibility

made it more practical to rely on published data than to implement or observe an AI tool in an actual hiring process for research purposes. Within this qualitative design, the approach was intentionally flexible and emergent. As is typical in qualitative inquiry, the research design was not a rigid blueprint set in advance but evolved as understanding deepened; for example, initial broad reading helped refine which specific ethical issues to focus on (Creswell and Poth, 2018, p. 47). However, the overall structure remained consistent with qualitative best practices: clearly defined research questions, systematic data collection via literature and document gathering, and a well-documented analytical procedure, thematic analysis.

It is important to note that this study does not fall neatly into one of the traditional qualitative approaches, such as ethnography, phenomenology, or grounded theory, because it does not involve fieldwork with participants or a bounded case as such. If one were to categorize it, it has elements of a case study, in the sense that it intensively examines a particular “case,” the case being the phenomenon of AI-based recruitment in the South African context, by drawing on multiple sources of evidence (Yin, 2018, p. 16). It is also akin to a qualitative content analysis of texts discussing that phenomenon (Schreier, 2012, p. 3). In designing the research this way, the intent was to remain focused on the research problem, ethical implications, and frameworks while leveraging existing knowledge. This design facilitates a contextualized understanding: it anchors the ethical analysis in real-world examples and discussions found in the literature, thereby ensuring that conclusions are not purely speculative but grounded in documented reality. In summary, the research design can be described as a qualitative desktop study that is exploratory and interpretive. It capitalizes on secondary sources to explore the research questions in depth. This design choice supports the study’s aims by allowing a broad examination of evidence from global studies to local regulations through a qualitative lens, building a rich narrative and thematic understanding of AI ethics in recruitment.

3.3 Desktop Research (Secondary Data)

Data collection for this study was conducted via desktop research, meaning the collection of secondary data from existing sources. This involved systematically identifying, gathering, and reviewing documents that could shed light on AI-based recruitment practices and their ethical implications. The types of sources consulted include academic journal articles and conference papers, industry and consulting reports, government and policy documents such as legislation or guidelines related to AI or employment in South Africa, case studies reported in the media or professional outlets, and relevant books or chapters. In qualitative terms, these documents and texts constitute the data set for analysis. Using documents as data has distinct advantages

for an ethical inquiry of this nature. It is an efficient and unobtrusive method many relevant documents are publicly available, especially through online databases, and can be obtained and analysed without affecting any subjects. Bowen (2009, p. 27) notes that “document analysis is a systematic procedure for reviewing or evaluating documents,” involving finding, selecting, and interpreting textual data. Following this guidance, a structured search and retrieval process was employed. Academic databases such as Scopus and Google Scholar were queried with keywords like “AI recruitment ethics,” “algorithmic hiring South Africa,” and “AI bias employment.”

Policy and legal databases were also searched for South African statutes or regulations concerning technology and hiring, for instance, checking the Protection of Personal Information Act (POPIA) for privacy aspects, or the Employment Equity Act (EEA) for relevance to algorithmic bias. The search further extended to institutional and professional bodies, such as the South African Human Rights Commission and the Department of Labour, which occasionally publish documents addressing fairness or discrimination in employment. Comparative insight was also gained from international frameworks like the EU General Data Protection Regulation (GDPR) and AI Ethics Guidelines from the European Commission (European Commission, 2019, p. 6). All gathered documents were catalogued, and their bibliographic details were recorded to ensure proper citation and reference management.

Inclusion criteria were established to maintain relevance: sources had to address at least one of the following domains: use of AI or algorithms in recruitment/hiring, ethical, legal, or social implications of AI in HR, or broader AI ethics frameworks applicable to recruitment. Preference was given to sources focusing on South Africa or similar developing/emerging economy contexts to ensure local relevance. However, foundational international works were also included to frame universal ethical issues, such as algorithmic bias and fairness (Ajunwa, 2020, p. 12). The period of literature considered was from 2010 onwards, with greater emphasis on the last 5–7 years, reflecting the period in which AI recruitment gained prominence. Earlier works were included only when foundational for example, core ethics theories or early discussions of automation in HR (Bentham, 1789, p. 15). Throughout the data collection, careful attention was paid to evaluating the credibility and perspective of sources. Peer-reviewed academic publications and official policy reports were granted higher evidential value, while media articles or blog posts were used only for illustrative cases and verified where possible.

This approach helped reduce the risk of relying on unsubstantiated claims and improved the reliability of insights drawn from secondary data (Creswell and Poth, 2018, p. 43). In total, a substantial corpus of documents was assembled, providing a multifaceted view of AI-based recruitment and its ethical challenges. As Bowen (2009, p. 29) highlights, qualitative document analysis requires that data be “examined and interpreted to elicit meaning, gain understanding, and develop empirical knowledge.” This interpretive process began during the data collection itself: as documents were reviewed, notes were made regarding key ethical issues raised, like bias, transparency, accountability, contextual factors, for example, South Africa’s legal environment, and theoretical insights like applications of consequentialist or deontological ethics. These notes later informed thematic coding and analysis, described in Section 3.5. In summary, this study employed desktop research as a qualitative document analysis method to gather secondary data. It provided a comprehensive and ethically non-intrusive means of collecting evidence. While this approach inherently depends on the availability and quality of existing documents, a limitation discussed later, it was well-suited to an exploratory ethics inquiry, allowing the researcher to draw upon collective scholarly and policy-based knowledge without the logistical and ethical complexities of participant-based research. All data were overseen with academic integrity: sources were accurately cited, and findings were interpreted faithfully within their original context.

3.4 Selection of Sources

Given that this study did not involve selecting human participants, the concept of sampling applies to the selection of documents and sources for analysis. A purposive sampling strategy was employed to deliberately select information-rich sources most relevant to the research questions. Purposive or purposeful sampling is a common technique in qualitative research wherein cases or data sources are selected based on their expected usefulness in illuminating the phenomenon of interest. As Patton (2015, p. 264) explains, the strength of purposive sampling lies in “selecting information-rich and relevant cases” that can yield an in-depth understanding. In this study, the “cases” were the publications and documents themselves, chosen intentionally rather than randomly to ensure their conceptual relevance to the ethical implications of AI-driven recruitment. The inclusion criteria were as follows. First, relevance to AI in recruitment: each selected source had to address AI, algorithms, or automated systems as used in recruitment, hiring, or related HR functions, or, in the case of broader AI ethics sources, contain sections applicable to hiring. Second, relevance to ethical implications: sources needed to discuss ethical issues, risks, benefits, or decision-making frameworks

directly, for example, studies on algorithmic bias in hiring. or indirectly, for example, legal analyses on data privacy relevant to recruitment data. Third, contextual relevance: priority was given to sources specific to South Africa, such as studies conducted locally, South African laws and codes like the King IV Report on Corporate Governance, which emphasizes ethical and responsible leadership, or examples from South African organizations adopting AI systems.

When directly relevant local sources were scarce, since AI in HR remains an emerging field, the sample was supplemented with international sources offering strong theoretical or practical insights. For instance, foundational works on algorithmic bias by Ajunwa (2020, p. 12) and global frameworks from the World Economic Forum (2025, p. 8) were included to contextualize how these issues may manifest within South Africa. Similarly, ethical perspectives from Bentham (1789, p. 15) and methodological frameworks from Creswell and Poth (2018, p. 50) and Braun and Clarke (2019, p. 593) were incorporated to support theoretical and methodological depth. Using these criteria, the researcher systematically screened search results and bibliographies of initial papers to select a robust and thematically diverse sample of documents. This purposive sample was not designed to be exhaustive of all literature on the topic; rather, it was sufficiently broad and varied to ensure comprehensive coverage of the key ethical and contextual dimensions of AI-based recruitment. The sample included peer-reviewed journal articles by AI and HR scholars for academic rigor, reports by global and South African organizations such as the World Economic Forum and the South African Human Rights Commission for industry and policy context, and local academic papers, for example, recent South African Journal of Human Resource Management studies offering empirical insights. Foundational methodological texts, such as Creswell and Poth (2018) and Braun and Clarke (2019), though not focused on AI, were also included because they informed the study's design and analysis framework.

In addition, legal and regulatory documents were purposively sampled to capture the governance aspect of AI recruitment. These included the Protection of Personal Information Act (POPIA, 2013) for data protection considerations, available Department of Labour guidelines on fair employment practices, and international governance frameworks such as the EU General Data Protection Regulation (GDPR) and the proposed EU AI Act for comparative insight (European Commission, 2019, p. 6). The sample thus incorporated a range of perspectives, technical, philosophical, ethical, and business-oriented, to strengthen the credibility of the findings through data triangulation. Since this purposive sampling method is non-probabilistic, the study makes no statistical claims of representativeness. Instead, it follows

the logic of saturation and sufficiency, where sampling continued until recurring themes indicated that new sources were no longer yielding novel insights. As Guest, Bunce, and Johnson (2006, p. 65) note, saturation occurs when additional data fail to generate new themes or information, signifying adequacy in coverage. The final corpus of sampled sources was thus determined to be sufficient for addressing the research questions. All selected sources were systematically documented, with notes on their relevance and inclusion rationale recorded for transparency and traceability. By employing purposive sampling, the study ensured that the data analysed were directly aligned with its ethical focus, thereby enhancing analytical depth and validity. This deliberate approach strengthens the reliability and richness of the ethical analysis undertaken in subsequent chapters.

3.5 Data Analysis: Thematic Analysis

The collected data were analysed using thematic analysis, a flexible qualitative method for identifying and interpreting patterns of meaning themes within data. Thematic analysis was chosen because it offers a systematic yet adaptable approach to organizing complex qualitative information, which is well-suited for an exploratory study involving diverse textual sources. According to Braun and Clarke (2006, p. 79), thematic analysis is “a method for identifying, analysing, and reporting patterns (themes) within data,” allowing researchers to both describe and interpret qualitative material. This study followed the six-phase process outlined by Braun and Clarke, namely (1) Familiarization with the data reading and re-reading documents while noting initial observations; (2) Generating initial codes systematically tagging significant segments of text (for example, passages about “bias in AI”); (3) Searching for themes collating codes into broader candidate themes for example, grouping codes related to fairness and discrimination under “Algorithmic Bias and Fairness”; (4) Reviewing themes, checking candidate themes against the dataset to ensure coherence and data support, refining or merging themes as necessary; (5) Defining and naming themes articulating each theme’s essence and relation to the research questions (for instance, defining what “Transparency and Accountability” means within AI recruitment; and (6) Producing the report weaving analytic narratives with evidence to present the findings. This process was adapted from Braun and Clarke’s (2006, pp. 86–93) foundational framework and refined in their later works on reflexive thematic analysis.

Throughout the analysis, the study employed a reflexive thematic approach, as articulated by Braun and Clarke (2019, p. 594), which recognises that researchers play an active role in knowledge production. In this view, themes do not simply “emerge” from data; rather, they are

the result of deep, iterative engagement between the researcher and the material. The researcher continuously reflected on how personal assumptions and theoretical orientation, particularly the ethical lens guiding the study, influenced coding and interpretation. This meant paying special attention to moral concepts such as fairness, accountability, and human dignity. Reflexivity, as noted by Creswell and Poth (2018, p. 82), enhances credibility by making visible the interpretive processes that shape meaning-making. Themes were not finalised until they had been examined from multiple angles and substantiated by several independent sources, a strategy that strengthened the dependability and trustworthiness of the findings. For example, in analysing a South African legal commentary on AI in hiring, the researcher noted phrases such as “no national AI regulatory framework” and “risk of indirect discrimination through algorithmic profiling.” These insights contributed to the synthesis of higher-level themes such as “Legal and Regulatory Gaps” and “Bias and Discrimination.” The analysis also incorporated a deductive or theoretical dimension, guided by the overarching ethical questions of the study. Braun and Clarke (2012, p. 58) affirm that thematic analysis can integrate both inductive and theoretical elements, provided the researcher is transparent about their influence. Thus, while codes emerged from the data, the ethical orientation of the study shaped the interpretive emphasis, particularly in identifying instances where fairness, transparency, or accountability were explicitly discussed.

During theme refinement, excerpts from multiple documents were collected to ensure each theme was well supported by textual evidence. For example, under the theme “Algorithmic Bias and Fairness,” extracts from diverse sources such as a corporate report acknowledging that AI may “inadvertently perpetuate biases” (Tech Financials, 2025, p. 4) and an academic discussion of bias in algorithmic tools, Ajunwa (2020, p. 13) were compared and synthesized. This triangulation within themes improved interpretive robustness by showing convergence and divergence across data types. In cases where contradictions arose, for example, one source claiming AI enhances objectivity while another emphasizes bias, the researcher explored contextual reasons for such differences rather than disregarding them, aligning with Patton’s (2015, p. 316) principle of analytic depth through contrast. In sum, the thematic analysis followed the rigorous procedures established by Braun and Clarke (2006, pp. 86-93), integrating inductive discovery with theoretical sensitivity. This ensured that findings were rooted in the data yet informed by ethical theory, yielding a nuanced understanding of how AI-driven recruitment intersects with ethical principles and regulatory realities in South Africa. The subsequent section elaborates on the DECA (Describe, Evaluate, Consult, Act) ethical

decision-making framework, which was used to deepen and structure the normative evaluation of these identified themes.

3.6 Methodological Framework: DECA (Describe, Evaluate, Consult, Act)

To complement the thematic analysis and explicitly tackle the ethical dimension of the research, the study incorporated the DECA framework as a methodological lens. DECA stands for Describe, Evaluate, Consult, Act, a four-step ethical decision-making model developed in an applied ethics context by scholars at the University of South Africa (Unisa) (Kretzschmar and Bentley, 2013). The DECA method is an adaptation of the classic “See–Judge–Act” approach to moral decision-making. It encourages a structured yet reflective analysis of ethical dilemmas by thoroughly describing the problem, evaluating options through ethical principles, consulting others or external standards, and acting based on informed judgment. Kretzschmar and Bentley (2013, p. 6) emphasize that ethical decision-making is not merely a top-down application of theory but a process that involves reflection on one’s decisions and openness to new moral questions. By using DECA, the researcher aimed to systematically assess each ethical issue identified in the thematic analysis and ensure that recommendations or conclusions were normatively sound and grounded.

3.6.1 Describe

The first step is to clearly describe the ethical issue or dilemma at hand (Kretzschmar and Bentley, 2013, p. 85). In this study, for each major theme emerging from the data, such as algorithmic bias, the nature of the issue was articulated in detail: What exactly is the ethical concern? Who is affected and how? What context surrounds it: technological, organisational, societal? The description phase involves being as objective as possible and mapping out the scenario or problem without rushing to judgment (Kretzschmar and Bentley, 2013, p. 5). For instance, when considering the use of AI CV-scanning tools that might disadvantage certain groups, the Describe step would outline the practice, note the observed outcomes, say, lower hiring rates for group X, and frame it as an ethical question of fairness in hiring. This ensured that each problem was explicitly spelled out and grounded in the data rather than assumed.

3.6.2 Evaluate

The second step of the DECA framework involves evaluating the ethical issue using established ethical theories and criteria. As outlined by Kretzschmar and Bentley (2013, p. 5), this step requires the application of multiple ethical lenses, notably deontological (duty-based),

teleological (consequence-based), and virtue ethics (character-based) perspectives, to each identified theme. This multidimensional approach enables a comprehensive ethical interpretation that accounts for rules, outcomes, and values. In practical terms, the researcher examined each ethical theme by asking Deontology (rules and rights): Are any rights being violated or duties unfulfilled by the AI practice? Is the practice consistent with moral principles or legal obligations? For example, does an AI hiring tool respect candidates' right to equal opportunity and privacy, or does it contravene principles enshrined in South African frameworks such as the Employment Equity Act (EEA) or Protection of Personal Information Act (POPIA)?, Teleology (consequences): What are the foreseeable outcomes of the AI practice, and do they produce a net benefit or harm? Who stands to gain or lose? For instance, does the use of AI improve efficiency and objectivity (positive outcomes) at the cost of excluding historically marginalised groups (negative outcomes)? This lens evaluates whether the overall impact of the technology aligns with the promotion of social welfare (Kretzschmar and Bentley, 2013, p. 5). Virtue Ethics (character and values): What would a morally exemplary organisation or HR professional do in this situation? Does the practice reflect virtues such as fairness, accountability, compassion, and integrity? This perspective considers whether the deployment of AI tools aligns with the moral character and ethical commitments that the organisation claims to uphold. By evaluating each theme through these three ethical lenses, the study illuminated the complex and layered nature of ethical issues in AI recruitment. For example, under deontology, algorithmic bias represents a violation of fairness and equality. Under teleology, it raises concerns about harmful consequences for excluded candidates, and under virtue ethics, it questions whether the organisation is acting with moral integrity and social responsibility. This triangulated ethical assessment ensured that the analysis moved beyond surface-level critique to a robust and principled evaluation of AI practices. It also reinforced the study's commitment to ethical pluralism, recognising that no single theory can fully capture the moral complexity of AI-driven recruitment in a post-apartheid South African context.

3.6.3 Consult

The third DECA step is Consult, which involves seeking external input or guidance. This stage recognizes that ethical dilemmas benefit from multiple perspectives and that no single viewpoint holds all moral truth (Kretzschmar and Bentley, 2013, p. 6). In this study, consultation involved engaging with wider literature and stakeholder perspectives, since there

were no live participants to consult directly. The researcher “consulted” academic scholarship, policy documents, and ethical codes to ensure the interpretations were supported by authoritative voices. For instance, stakeholder perspectives included regulators’ concerns about AI oversight, candidates’ anxieties about algorithmic bias, and developers’ justifications for AI tools. Ethical codes such as the South African Code of HR Ethics, King IV Report on Corporate Governance, and global standards like the IEEE Ethically Aligned Design and EU AI Guidelines were also consulted to contextualize moral obligations. This cross-referencing validated the study’s interpretations and reinforced its normative grounding.

3.6.4 Act

The last step, Act, entails making an informed ethical decision and implementing it. In academic research, “Act” corresponds to drawing conclusions and formulating practical recommendations. After describing, evaluating, and consulting, the researcher must determine a justified course of action (Kretzschmar and Bentley, 2013, p. 6). In this study, this took the form of proposing measures such as bias audits for AI systems, clearer regulatory guidelines, and stronger ethical training for HR professionals using AI. DECA also recognizes that ethical action is iterative new challenges may arise from implemented solutions. Thus, the study’s conclusions not only recommend ethical practices but also identify areas for further inquiry and continuous reflection. By integrating the DECA framework into the methodology, the research bridges the gap between identifying ethical issues via thematic analysis and normatively evaluating them. Kretzschmar and Bentley’s approach insists that ethics students must apply and embody ethical reasoning, not merely describe it. In this study, DECA provided a scaffold for applying ethical theory to practice. For example, when a theme like “lack of transparency” emerged, DECA prompted the researcher to describe the phenomenon, evaluate its ethical implications, consult external guidance, and act by recommending transparency-enhancing measures (Kretzschmar and Bentley, 2013, p. 6). The use of DECA as a methodological framework is a distinguishing feature of this study. It ensures that the methodology is not only qualitative but explicitly ethics oriented. This aligns perfectly with the research objective of examining ethical implications. It adds a layer of normative analysis to the descriptive thematic results. Moreover, DECA’s origin in a South African applied ethics context, as per Kretzschmar and Bentley (2013, pp. 1-9), makes it particularly apt, grounding the methodology in local scholarly practice. In summary, DECA provided a clear, stepwise

process to analyse and respond to ethical issues, enhancing the rigor and depth of the study's evaluation of AI-based recruitment in South Africa.

3.7 Conclusion

In conclusion, by addressing credibility, transferability, dependability, and confirmability, this research has taken comprehensive steps to ensure its trustworthiness. The use of multiple sources and methods, method triangulation by combining thematic analysis with an ethical framework, further adds to the robustness. The reader can thus have confidence that the study's insights into AI-based recruitment ethics are not only well-founded in evidence but also produced via a rigorous and transparent research process. This strong methodological foundation is particularly important given the interdisciplinary and evaluative nature of the topic, where both factual and value-based interpretations must be managed carefully. Having outlined the methodological structure and established the trustworthiness of this study, the next chapter presents the theoretical framework that underpins the analysis. The framework provides the philosophical and ethical lens through which the findings are interpreted. Specifically, this study adopts a consequentialist ethical framework with particular attention to utilitarian principles to evaluate the moral implications of using Artificial Intelligence in recruitment and selection within the South African context. This theoretical orientation helps in assessing whether AI recruitment systems promote overall well-being and fairness or, conversely, exacerbate inequality and exclusion. The following chapter thus explores the key tenets of consequentialism and related ethical theories, explaining how they guide the interpretation of the data and the ethical evaluation of AI-driven hiring practices.

CHAPTER FOUR: THEORETICAL FRAMEWORKS – THE ETHICAL THEORY OF CONSEQUENTIALISM

4.0 Introduction

Building on the literature review presented in Chapter 2, which traced the evolution of recruitment from manual screening to algorithmic shortlisting. This chapter synthesizes empirical findings on AI's effects on efficiency and bias, highlighting the unique challenges of South Africa's high-unemployment, diversity-driven context and drawing on the research methodology outlined in Chapter 3. This chapter turns to the theoretical lenses that will structure our analysis. Centred on a consequentialist perspective, I will interrogate AI-driven hiring by asking whether its outcomes truly maximize overall welfare while minimizing harm. Three contemporary scholars guide this inquiry: Peter Singer's work on effective altruism (2015) reinforces the consequentialist imperative to maximize well-being and reduce suffering in practical, evidence-based ways. Shlomi Segall (2021) offers a justice-oriented consequentialist lens, urging ethical scrutiny of structural and institutional effects rather than isolated actions. Andrés Mejía (2019) advances institution-sensitive consequentialism, highlighting the importance of evaluating ethical outcomes within specific socio-political and historical contexts, which is particularly relevant to post-apartheid South Africa. By situating these voices within a South African setting, a full-bodied conceptual foundation will be established for evaluating AI recruitment is promises and risks.

4.1 The ethical theory of Consequentialism

Consequentialism is a regulating ethical theory that evaluates the morality of actions based solely on their outcomes. An action is right, according to this framework, if it leads to the best overall results, typically defined in terms of welfare, utility, or well-being (Driver, pp. 97–98). Rather than appealing to fixed moral rules or intrinsic virtues, consequentialism asks: What will the consequences of this action be? It directs agents to choose the path that maximizes benefit and minimizes harm for all affected. Contemporary defenders of consequentialism have extended and refined the theory to address complex ethical challenges in a globalised, digitised world. de Lazari-Radek and Singer (2017) argue for a version of hedonistic utilitarianism that maintains Benthamite roots but stresses the importance of applying utilitarian principles pragmatically to real-world issues such as poverty, digital equity, and labour rights (de Lazari-Radek and Singer, 2017, pp. 62–65). Their advocacy of effective altruism has helped shape how ethical consequences are evaluated in global and technological contexts, including employment technologies like AI recruitment.

Mejía (2019) introduces a more contextual approach with his theory of institution-sensitive consequentialism, which argues that we cannot assess outcomes without considering the institutional frameworks in which actions are embedded. According to Mejía, moral evaluation must account for how systems such as employment policies, recruitment algorithms, and legal norms shape consequences over time (Mejía, 2019, p. 312–314). This is especially salient in South Africa, where the historical and structural context of inequality must inform ethical judgments about hiring practices. Further advancing the consequentialist paradigm, Segall (2021) offers a justice-oriented perspective by examining the long-term impacts of institutional design and structural disadvantages. He argues that systematic disparities in outcome, even when not caused by individual intention, should be morally scrutinised under consequentialist terms (Segall, 2021, p. 133). For AI recruitment tools, this implies that even apparently neutral algorithms can be ethically problematic if they reproduce social and economic marginalisation. Arrhenius, Bykvist, and Campbell (2020) focus on population ethics and the moral weight of future consequences, introducing tools to evaluate long-term and intergenerational outcomes. Their work is particularly relevant in assessing whether technologies like AI in recruitment will lead to sustained inclusion or entrenched exclusion over time (Arrhenius, Bykvist and Campbell, 2020, p. 115–120). This adds a forward-looking dimension to consequentialist ethics in technology governance.

Figure 1: The Ethical Theory of Consequentialism and Research Objectives for Evaluating AI in Recruitment



Own source (2025)

The diagram is a structured, flowchart-style visual representation that illustrates the research focus on evaluating the ethical implications of integrating Artificial Intelligence (AI) into job recruitment and selection processes in South Africa. It is organised in a clear top-down layout that guides the viewer from the central research aim to the specific sub-objectives and finally to the ethical framework underpinning the study.

At the very top of the diagram, a large dark-blue rectangular box contains the main research objective written in bold white text: *“To evaluate the ethical implications of integrating AI into job recruitment and selection processes in South Africa.”* This serves as the core focus from which all other components of the diagram flow.

Directly beneath this heading is a rounded white box labelled *“1.6 Research Sub-Objectives.”* This section acts as a bridge between the overarching aim and the three sub-objectives listed below. From this box, three connecting lines extend downward to three separate coloured blocks, each representing one sub-objective.

The first sub-objective, on the left, is displayed in an orange block. It reads: *“To analyse traditional recruitment processes in South Africa.”* This section highlights the importance of understanding conventional hiring practices before comparing them to AI-driven alternatives.

The second block, in the centre, is light blue and states: *“To investigate AI-driven recruitment mechanisms and challenges.”* This reflects the study’s interest in exploring how AI operates within recruitment, including its benefits and limitations.

The third block, also light blue, is positioned on the right and reads: *“To assess the ethical consequences of AI adoption in recruitment.”* This emphasises concerns such as fairness, bias, transparency, and accountability.

Finally, at the bottom of the diagram, a dark-blue block contains the phrase *“Ethical Theory of Consequentialism.”* This suggests that the entire evaluation is grounded in consequentialism, an ethical framework that evaluates actions based on their consequences.

4.2 Definition and meaning of Consequentialism

Consequentialism is an ethical theory that evaluates the morality of actions based on their outcomes or consequences rather than intentions, motives, or intrinsic nature. According to Bentham (1789, p. 14), “it is the greatest happiness of the greatest number that is the measure

of right and wrong.” This principle means that an action is morally right if it leads to the best possible overall results for the majority. Similarly, Mill (1861, p. 9) defines consequentialism as a moral doctrine where the rightness of an action depends on the amount of happiness or utility it produces. Consequentialism determines moral value by assessing outcomes. A decision is good if it maximizes benefits and minimizes harms (Sinnott-Armstrong, 2022, p. 3). This outcome-focused reasoning stands in contrast to deontological ethics, which focuses on duties and rules, and virtue ethics, which emphasizes character traits (Arrhenius, Bykvist, and Campbell, 2020, p. 115). Philosophically, consequentialism provides a rational framework for moral evaluation in social systems, including modern technologies like Artificial Intelligence (AI)(Limarga, Sutoyo, and Santoso, 2023). When applied to recruitment and selection, consequentialism asks: Do AI systems produce fairer, more beneficial outcomes for society and job seekers? For instance, if AI tools reduce human bias and speed up hiring, the overall consequences might be positive (Floridi and Cows, 2021, p. 697). Conversely, if algorithms exclude marginalized applicants, the negative outcomes render the system unethical (Ajunwa, 2020, p. 12). From a South African perspective, consequentialism aligns with the Constitution’s emphasis on equity and social welfare. Ethical evaluation must therefore consider not only business efficiency but also broader social outcomes, such as inclusion and justice (Mkhize, 2021, p. 45). In recruitment, the moral test is whether AI contributes to equitable employment or deepens inequality. Thus, consequentialism offers a pragmatic ethical lens through which the effects of technological innovation can be assessed, considering South Africa’s socio-economic realities.

4.3 Approaches of Consequentialism

There are several approaches within consequentialism, each offering a unique way to determine and measure “good” outcomes. The three dominant approaches are act consequentialism, rule consequentialism, and motive consequentialism.

4.3.1 Act Consequentialism

Act consequentialism judges the morality of each specific action based on its direct consequences (Smart and Williams, 1973, p. 75). In this view, the right act is the one that, among all available options, produces the best outcomes. Bentham’s utilitarianism exemplifies this form; each decision must be analysed by its utility, pleasure over pain for those affected (Bentham, 1789, p. 22). In AI recruitment, act consequentialism might assess whether using an algorithm in a specific hiring instance leads to fairer outcomes. For example, if an AI screening

tool reduces discriminatory decisions in one recruitment process, the use of its tool is justified. However, if the same system unfairly rejects qualified candidates, the act becomes unethical.

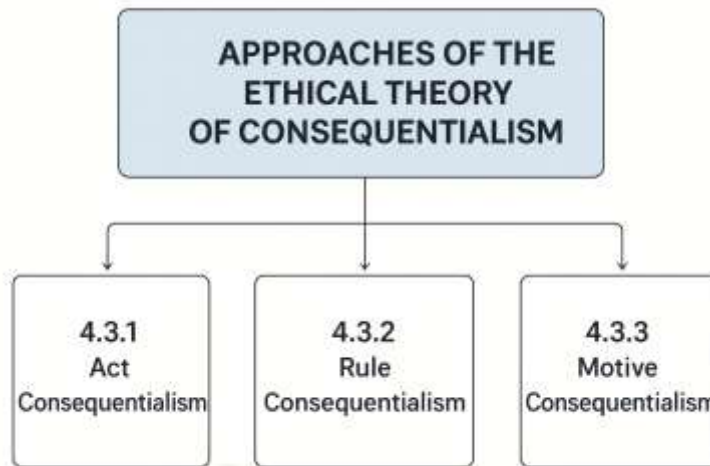
4.3.2 Rule Consequentialism

Rule consequentialism, advanced by Hooker (2000, p. 32), evaluates morality based on adherence to rules that, if followed, produce the best results. Instead of judging each act individually, this approach looks at the long-term consequences of adopting consistent ethical norms. In recruitment ethics, a rule consequentialist might argue that implementing transparent AI policies, such as fairness audits or accountability protocols, produces the best overall outcomes for society. Such rules promote trust, protect human rights, and mitigate harm caused by biased algorithms. This aligns with South Africa's Employment Equity Act (No. 55 of 1998), which sets fairness rules for hiring to achieve equitable outcomes.

4.3.3 Motive Consequentialism

This approach considers the motives behind actions as far as they produce desirable outcomes (Adams, 1976, p. 471). An act may be morally right if motivated by intentions that tend to generate satisfactory results, even if an individual instance fails. For example, a company introducing AI recruitment with the motive of improving diversity might be ethically justified, provided the overall intention promotes fairness and inclusion. In the South African context, consequentialist motives underpin transformative justice, as employers adopt AI for social benefit, reduce bias, widen access, and fulfil moral responsibility despite technical imperfections (Sinnott-Armstrong, 2022, p. 6). Overall, these approaches demonstrate the flexibility of consequentialism as a moral framework; whether assessing individual acts, general rules, or guiding motives, the ethical criterion remains the consequences of decisions for collective well-being.

Figure 2: Representing approaches of the ethical theory of consequentialism



Own source (2025)

The diagram illustrates the three main approaches of consequentialism: Act, Rule, and Motive Consequentialism. Act Consequentialism focuses on judging each action by its individual outcomes, whether it produces the greatest good for the greatest number. Rule Consequentialism evaluates the consequences of following general moral rules, emphasising consistency and fairness in ethical decision-making. Motive Consequentialism, on the other hand, considers the moral worth of actions based on the intentions or motives behind them, as well as their effects.

This framework is important when critiquing the ethical effects of using AI in job recruitment and selection in South Africa, as it helps assess decisions and systems not only by their immediate outcomes but also by their adherence to ethical rules and the intentions behind their design. Applying these approaches ensures a balanced ethical evaluation of addressing fairness, transparency, discrimination, and the broader societal impact of automated hiring systems.

4.4 Different Types of Consequentialism

Consequentialism branches into several types depending on what is considered the ultimate “good” outcome. The main types include utilitarianism, hedonistic consequentialism, preference consequentialism, ethical egoism, and negative consequentialism.

4.4.1 Utilitarianism

The most well-known form, utilitarianism, was pioneered by Jeremy Bentham and John Stuart Mill. It asserts that moral actions maximize happiness or utility for the greatest number (Bentham, 1789, p. 15; Mill, 1861, p. 10). In AI recruitment, utilitarianism supports adopting

technologies that improve efficiency and fairness for both employers and job seekers (Binns, 2018, p. 3). However, if such systems cause widespread unemployment or discrimination, they fail the utilitarian test of producing the greatest happiness.

4.4.2 Hedonistic Consequentialism

This version identifies pleasure or happiness as the only intrinsic good (Crisp, 1997, p. 12). Actions are right if they increase pleasure or reduce pain. In recruitment ethics, hedonistic consequentialism could imply that AI is moral if it reduces stress and improves satisfaction for job seekers and recruiters alike. Yet, if it introduces anxiety or alienation, for example, applicants feeling dehumanized by automated systems, the negative psychological effects outweigh the benefits.

4.4.3 Preference Consequentialism

Proposed by R. M. Hare (1981, p. 92) and further developed by Peter Singer (1993, p. 10), this type focuses on satisfying individuals' informed preferences rather than mere pleasure. It considers actions right if they fulfil people's rational desires. In recruitment, this view supports AI tools that respect candidate autonomy and fairness, such as allowing applicants to appeal algorithmic decisions or explain the data used.

4.4.4 Ethical Egoism (Individual Consequentialism)

Ethical egoism claims that actions are morally right if they promote one's self-interest (Rand, 1964, p. 25). While not widely accepted as a social ethic, in business contexts, it emerges when organizations justify the use of technology solely for profit or efficiency gains. In South Africa, such egoistic consequentialism may conflict with collective ethics rooted in Ubuntu and social justice, rendering it ethically limited in addressing employment inequality.

4.4.5 Negative Consequentialism

This type prioritizes reducing harm or suffering rather than maximizing pleasure (Popper, 1945, p. 151). The ethical goal is to minimize negative outcomes. In AI recruitment, this principle suggests that systems should be designed primarily to avoid harm, such as eliminating discriminatory patterns or preventing job loss due to automation. Floridi and Cowls (2021, p. 698) note that negative consequentialism aligns with global AI ethics principles, emphasizing "do no harm" as a baseline.

4.4.6 Ideal Consequentialism

Moore (1903, p. 83) proposed that other values, such as beauty, knowledge, and friendship, also have intrinsic worth. Thus, actions are right if they promote a combination of these goods. Applying this to AI recruitment means valuing not only efficiency but also human dignity, trust, and social inclusion.

4.5 Proponents of Consequentialism

Jeremy Bentham (1789) and John Stuart Mill (1879) are the classical architects of consequentialist ethics, particularly the utilitarian tradition. Bentham's principle of utility, often summarized as "the greatest happiness for the greatest number," established that the moral rightness of an action depends on the extent to which it increases overall pleasure or reduces pain for all affected. Mill later refined utilitarianism by emphasizing the quality of pleasures, arguing that intellectual and moral pleasures hold higher value than purely physical ones (Mill, 1879, p.96). Together, Bentham and Mill laid the foundation for evaluating morality based on aggregate welfare, a foundation upon which modern consequentialist theorists continue to build.

Contemporary consequentialist thought has evolved significantly, with modern scholars applying its principles to urgent technological, social, and political challenges. In the context of AI-driven recruitment and selection, recent proponents of consequentialism provide valuable tools for evaluating how outcomes, both immediate and long-term, shape the ethical permissibility of technological interventions in the labour market. Driver (2015, pp. 97–100) reinvigorates consequentialist ethics by emphasizing the importance of practical reasoning in moral judgment. She argues that consequentialism's adaptability to empirical evidence makes it particularly well-suited to applied ethics in technologically evolving contexts. Her work critiques rigid ethical systems that fail to account for outcome-based analysis in complex social environments like hiring algorithms. This adaptability is especially relevant when assessing the fairness and societal impact of AI recruitment tools, which must be evaluated by looking at actual results such as patterns of inclusion or exclusion across demographic groups.

Mejía (2019, pp. 312–314) proposes a nuanced version of consequentialism called "institution-sensitive consequentialism," which considers the broader institutional structures in which decisions occur. He contends that ethical choices, especially those involving algorithmic systems, cannot be divorced from the socio-political contexts they operate in. For South Africa, a country with a legacy of structural inequality, this perspective is especially pertinent. Mejía's

framework encourages assessing AI recruitment technologies not just for their efficiency, but for how they may reinforce or disrupt institutional patterns of Proponents of consequentialism argue that it provides a transparent, rational, and impartial basis for ethical decision-making by linking moral judgments to the outcomes of actions. This approach enables comparability and universality because consequences can be evaluated fairly across different situations. Bentham (1789, pp. 2–4) claimed that morality should be guided by the principle of utility, which is the tendency of an action to produce the greatest happiness for the greatest number.

Mill (1879, pp. 10–12) improved this utilitarian approach by distinguishing between higher and lower pleasures, asserting that intellectual and moral satisfactions have greater intrinsic value than physical pleasure alone. Together, these classical thinkers laid the groundwork for consequentialist reasoning as a moral framework focused on maximizing welfare and ensuring impartiality. Arrhenius, Bykvist, and Campbell (2020, pp. 115–120) have contributed to consequentialist population ethics, developing models that consider how policies and technologies affect not only current populations but also future generations. This long-view analysis is crucial in evaluating AI's impact on intergenerational equity, particularly in developing economies.

de Lazari-Radek and Singer (2017, pp. 62–65) extend traditional utilitarianism to contemporary issues, including global justice, digital ethics, and effective altruism. Their collaborative work defends a hedonistic utilitarian framework while advocating practical applications that minimize suffering and maximize well-being, especially for the worst-off. In the context of AI hiring, this implies a moral imperative to design systems that actively improve access to employment for historically disadvantaged communities rather than merely increasing corporate efficiency. Segall (2021, p. 133) provides a consequentialist critique of structural injustice, arguing that institutions bear responsibility not just for individual actions but for sustained inequalities in outcomes. This position is particularly relevant when evaluating AI-driven recruitment: even if no human decision-maker intentionally discriminates, an algorithmic system can produce unethical consequences if it systematically disadvantages certain groups. Such insights challenge AI developers and employers to proactively mitigate algorithmic bias, thereby avoiding the perpetuation of unjust employment outcomes. Together, these modern scholars shift consequentialism from abstract theory to ethical praxis, emphasizing not only what actions intend to achieve but also what they produce in real-world scenarios.

4.6 Arguments for the Ethical Theory of Consequentialism

Proponents of consequentialism believe it provides a clear, logical, and unbiased basis for ethical decision-making by linking moral judgments to the outcomes of actions. This approach enables comparability and universality because consequences can be evaluated fairly across various situations. Bentham (1789, pp. 2–4) argued that morality should be guided by the principle of utility, which is the tendency of an action to produce the greatest happiness for the greatest number. Mill (1879, pp. 10–12) refined this utilitarian approach by distinguishing between higher and lower pleasures, arguing that intellectual and moral satisfactions are more valuable than simple physical pleasures. Together, these classical thinkers laid the groundwork for consequentialist reasoning as a moral system focused on maximizing welfare and maintaining fairness.

Contemporary theorists have emphasised the adaptability and empirical grounding of consequentialism. Driver (2015, pp. 97–100) highlights that consequentialism accommodates evolving moral and social conditions by incorporating new empirical evidence into ethical reasoning. This dynamic quality allows the theory to respond flexibly to the complexities of modern society. Miklós (2025, p. 211) expands this notion through the idea of institutional consequentialism, which argues that moral burdens can be shared across institutions rather than placed solely on individuals, making the theory less demanding while preserving its focus on collective outcomes. Likewise, Andersen (2024, pp. 84–86) broadens consequentialist analysis into the domain of epistemic consequentialism, maintaining that even beliefs and judgments should be evaluated by their consequences for knowledge and well-being.

Further contributions by Wyner and Stachura-Zurek (2024, pp. 33–37) propose value-based consequentialism, integrating personal values into outcome assessments to mitigate concerns about moral alienation. These developments demonstrate that modern consequentialism retains its relevance and rigor while accommodating moral pluralism. Consequently, the theory continues to serve as a systematic and outcome-oriented ethical model capable of guiding moral agents, policymakers, and institutions toward actions that promote the greatest aggregate welfare.

Building on these key points, consequentialism's strength lies in its clarity, universality, and practical use in moral decision-making. By emphasizing outcomes rather than intentions or duties, consequentialism provides an ethically measurable framework for judging right and wrong. Unlike deontological theories that focus on moral rules regardless of their consequences, consequentialism states that the moral value of an action depends on its actual

effects (Driver, 2015, pp. 97–100). This focus on results makes it especially useful for evaluating technologies like Artificial Intelligence, where the impacts, rather than the intentions, determine the moral worth of innovation. Additionally, consequentialism's empirical flexibility allows it to adapt as scientific and societal developments occur. As Driver (2015, pp. 102–103) elaborates, consequentialist reasoning evolves by incorporating new data into moral evaluations, thereby keeping ethical judgments relevant to changing conditions. This adaptability is especially important in AI ethics, where consequences such as algorithmic bias, data privacy breaches, or increased hiring efficiency can be measured and compared. Similarly, Miklós (2025, p. 211) introduces the idea of institutional consequentialism, proposing that moral responsibility can be shared across systems and organizations instead of being assigned only to individuals. In AI recruitment, this view supports collective accountability, where both developers and employers share the duty to ensure that algorithmic systems promote social welfare and reduce harm.

Andersen (2024, pp. 84–86) extends the reach of consequentialism to epistemic contexts, proposing that beliefs and judgments should also be evaluated by their outcomes for knowledge, understanding, and human well-being. This form of epistemic consequentialism reinforces the idea that moral assessment should consider how actions contribute to informed decision-making and truth-seeking principles that are crucial for ethical AI design, where transparency and accuracy are key determinants of trust. Further contributions by Wyner and Stachura-Zurek (2024, pp. 33–37) introduce value-based consequentialism, which integrates personal and cultural values into the evaluation of consequences. This approach acknowledges that while the theory prioritizes aggregate welfare, it must also respect moral diversity and contextual values. By allowing local ethical frameworks such as South Africa's ubuntu ethics to inform consequence evaluation, consequentialism becomes more inclusive and less mechanistic in application. Taken together, these developments demonstrate that modern consequentialism retains both relevance and theoretical rigor. It has evolved from a rigid utilitarian calculus into a flexible, empirically informed, and value-sensitive ethical model. As such, consequentialism provides a coherent foundation for evaluating the ethical implications of AI in recruitment and selection. Its focus on outcomes aligns with promoting fairness, efficiency, and overall social welfare in organizational contexts. As Miklós (2025, p. 212) notes, consequently, the theory continues to serve as a systematic, outcome-oriented moral framework capable of guiding moral agents, policymakers, and institutions toward actions that promote the greatest aggregate good.

4.7 Arguments Against the Ethical Theory of Consequentialism

Despite its theoretical appeal, consequentialism faces enduring criticisms that question its philosophical coherence and moral sufficiency. The attractiveness of the theory lies in its apparent rational simplicity, judging moral worth solely by outcomes, but this very simplicity has generated complex philosophical concerns about how it treats persons, duties, and moral knowledge. Critics argue that in attempting to construct a universal calculus of good, consequentialism risks ignoring the qualitative dimensions of morality that cannot be reduced to numerical or utilitarian evaluation. This has prompted an extensive scholarly debate regarding whether a framework so focused on results can adequately protect the moral worth of individuals or sustain the integrity of ethical decision-making in diverse social contexts.

A major objection concerns the separateness of persons, the claim that consequentialism fails to respect individual dignity by aggregating distinct lives into a collective calculus. Fischer (2024, pp. 59–61) contends that such aggregation overlooks the moral significance of personal boundaries, effectively reducing individuals to interchangeable units of welfare. This problem, originally raised by Rawls and later echoed by modern critics, challenges the very core of consequentialism’s utilitarian foundations. By collapsing individual experiences into a shared moral total, the theory risks eroding the inviolability of individual rights. In this sense, consequentialism may allow morally troubling trade-offs, for instance, justifying harm to one person if it benefits many others. This critique suggests that a purely consequentialist ethic may permit sacrificing one person’s well-being for the benefit of many, undermining justice, and equality. Within applied ethics, such as in AI-driven recruitment systems, this objection becomes strikingly relevant: if an algorithm systematically disadvantages certain applicants but overall improves institutional efficiency, consequentialism might wrongly deem this “morally acceptable” under the principle of utility maximization.

Another recurring challenge is the over-demandingness objection, which holds that consequentialism imposes unrealistic moral obligations on agents to always act in ways that maximize overall utility. This criticism highlights the impracticality of expecting moral agents, particularly ordinary individuals, to continually calculate the best possible outcome in every circumstance. Moral life, critics argue, is filled with competing responsibilities, emotional constraints, and limited resources; requiring constant utility maximization ignores these human realities. Miklós (2025, p. 214) responds by arguing that moral duties should be distributed between individuals and institutions, yet this approach remains contested, particularly in societies where institutional capacity is limited. In such contexts, the burden of moral action

often falls on individuals who lack the power or resources to effect large-scale good, highlighting a gap between ethical theory and social reality. The over-demandingness critique therefore reveals consequentialism's potential alienation from moral psychology it risks transforming ethics into an abstract exercise of moral accounting rather than a guide for human flourishing.

Klement (2023, pp. 91–94) further argues that consequentialism neglects the moral significance of motives, integrity, and virtue, reducing moral worth to external outcomes while disregarding internal character. This critique speaks to a deeper philosophical division between agent-centred and outcome-centred moralities. While consequentialism values results, virtue ethics, and deontology value the moral agent and their intentions. By ignoring internal moral dispositions, consequentialism may fail to explain why acts performed from good motives are often judged as morally superior, even if their results are imperfect. This critique aligns with Bennett's (2022, p. 128) observation that consequentialist theories of punishment risk treating individuals merely as means to social ends, ignoring relational justice and moral responsibility. In institutional ethics, such as organizational hiring decisions, this could translate into valuing outcomes like productivity or efficiency over fairness, empathy, or compassion. The risk is that moral reasoning becomes instrumental, losing sight of the human and relational dimensions of ethical life.

Finally, scholars such as Stenseke (2023, pp. 144–147) emphasize practical and computational limits in applying consequentialist reasoning. Because real-world agents often lack complete information or predictive ability, the precise calculation of outcomes is frequently infeasible. This epistemic limitation calls into question whether consequentialism can serve as a practical guide to action rather than merely a theoretical ideal. Even with the advent of sophisticated data analytics and AI, the moral consequences of decisions often unfold unpredictably over time, defying precise quantification. As a result, consequentialism may provide elegant moral logic but fail to deliver workable moral guidance in uncertain, complex environments such as technology governance or employment ethics.

Collectively, these objections highlight the tension between consequentialism's conceptual clarity and its potential moral shortcomings, suggesting the need for hybrid frameworks that balance outcome-based reasoning with respect for rights, virtues, and human agency. This conclusion has inspired the rise of integrative ethical approaches such as rule consequentialism, contractual, and virtue-based consequentialism, which attempt to reconcile the strengths of

consequence-oriented ethics with safeguards for justice and autonomy. In the context of AI recruitment ethics, adopting such hybrid models enables a more holistic evaluation of technological practices, ensuring that efficiency and innovation do not override moral accountability or human dignity.

4.8 The Role of The Ethical Theory of Consequentialism in Guiding the Work

This study adopts consequentialism as its guiding ethical framework to evaluate the legitimacy of AI-driven recruitment practices in South Africa. Consequentialism, which assesses the morality of actions based on their outcomes, provides a structured lens through which the benefits and harms of algorithmic hiring can be weighed. Rather than focusing on technological novelty or procedural compliance alone, this approach centres on the real-world effects of AI systems on candidates, employers, and society. In applying this theory, the study will examine whether AI recruitment tools contribute to collective welfare, fairness, and social inclusion, especially within South Africa's post-apartheid socio-legal context. Consequentialism enables a holistic assessment that considers both quantitative indicators, for example, efficiency metrics, diversity outcomes, and qualitative dimensions, for example, dignity, trust, and lived experience.

It also supports a justice-sensitive adaptation, aligning ethical evaluation with local imperatives such as the Employment Equity Act and Broad-Based Black Economic Empowerment policies. The theory will be used to interpret thematic findings from the data, guiding the ethical analysis in Chapter Five. Each identified issue, such as bias, opacity, or exclusion, will be assessed in terms of its overall impact on well-being, rather than its technical or procedural characteristics. This outcome-oriented methodology ensures that the study remains focused on both ethical legitimacy and operational functionality. By framing the inquiry through consequentialism, the research maintains a normative commitment to evaluating whether AI recruitment systems advance or undermine social justice, and whether they produce outcomes that are morally defensible in the South African context.

4.9 Conclusion

This chapter places the ethical examination of AI-driven recruitment within a consequentialist framework, demonstrating how outcome-oriented reasoning clarifies the technology's dual potential: to enhance efficiency, fairness, and data-driven precision, or to reinforce opacity and dehumanize the hiring process. I traced the evolution of consequentialist thought from Bentham and Mill's classical utilitarian calculus to the modern perspectives of Singer, Segall, Driver,

and Mejía, emphasizing the relevance of this theory in assessing AI's practical impacts in South Africa. In the previous chapter (Chapter 3), the research methodology used a desktop approach with the DECA method to systematically map existing knowledge. This evidence base, combined with the theoretical framework outlined here, provides a solid foundation for the upcoming empirical investigation. The next chapter, Chapter 5: Analysis, will apply the established consequentialist framework to explore the ethical implications of AI-driven recruitment and selection. Using a thematic analysis of secondary data, it will examine how AI systems in South African hiring practices conform to or diverge from consequentialist principles of maximizing overall welfare, fairness, and social good. Hence, Chapter 5 will translate theoretical reasoning into practical ethical evaluation, assessing the real-world consequences of AI technologies within both national and global employment contexts.

CHAPTER FIVE: ANALYSIS OF AI RECRUITMENT WITH THE ETHICAL THEORY OF CONSEQUENTIALISM

5.0 Introduction

This chapter presents a comprehensive ethical analysis of Artificial Intelligence (AI) in recruitment and selection, using the moral theory of consequentialism as its guiding framework. Driver (2015, p. 19) observes that consequentialism holds that the ethical rightness or wrongness of an action is determined solely by its outcomes, specifically, whether it produces more good than harm for all affected stakeholders. This outcome-oriented lens is particularly well-suited for evaluating AI-driven recruitment systems, which are reshaping the employment landscape in ways that have profound implications for fairness, justice, and social welfare. The global integration of AI into human resource management has transformed traditional recruitment paradigms.

Employers now use AI to streamline hiring processes, identify talent more efficiently, and reduce administrative burdens. Chamorro-Premuzic et al (2020, p. 78) argue that these technologies promise significant benefits, including enhanced speed, objectivity, and predictive accuracy in candidate evaluation. However, these advancements also introduce complex ethical dilemmas, particularly around issues of transparency, accountability, algorithmic bias, and respect for human dignity. In practice, the consequences of AI implementation in recruitment extend far beyond operational efficiency. They directly affect the livelihoods, opportunities, and psychological well-being of job seekers. As such, AI recruitment must be examined through a moral lens that prioritises collective welfare, equity, and justice, especially in contexts marked by historical and structural inequality.

From a consequentialist perspective, the ethical legitimacy of AI recruitment does not rest on the intentions of developers or the technical sophistication of algorithms, but rather on the real-world outcomes these systems produce. If AI tools demonstrably improve fairness, accelerate access to employment, and reduce discriminatory practices, they can be considered morally valuable because they enhance collective well-being. Binns (2020, p. 520) argues that, conversely, if they entrench bias, obscure accountability, or diminish human oversight, they generate harmful consequences that undermine their moral legitimacy. This ethical calculus is particularly urgent in South Africa, where the labour market continues to reflect the enduring legacies of apartheid, including persistent racial, gender, and class disparities mentioned by Nkosi and Ruggunan (2023, p. 18). The Employment Equity Act 55 of 1998 and Section 9 of the Constitution of the Republic of South Africa (1996) impose both legal and moral

obligations on employers to eliminate unfair discrimination and promote equitable representation. In this context, the deployment of AI in recruitment cannot be ethically neutral or technologically deterministic. It must be deliberately aligned with the country's broader goals of redress, inclusion, and transformation.

When assessed through the consequentialism framework, AI recruitment must therefore demonstrate that its aggregate outcomes, both immediate and long-term, contribute positively to the moral and socio-economic welfare of South Africa's workforce. This includes not only improving hiring efficiency but also advancing substantive equality, enhancing access to opportunity, and reinforcing institutional trust. This chapter proceeds by applying consequentialist reasoning to weigh the ethical benefits and risks associated with AI recruitment. It begins by exploring the positive consequences of AI adoption, including increased efficiency, reduced human bias, and cost savings outcomes that are morally desirable under a utilitarian perspective that prioritises the maximisation of collective good.

The discussion then turns to the ethical risks and unintended consequences of AI in recruitment, such as algorithmic discrimination, lack of transparency, dehumanisation, and over-reliance on automated systems. Each of these concerns is examined in terms of its impact on fairness, human dignity, and social justice. Through this analysis, the chapter seeks to determine whether the use of AI in recruitment within South Africa can be ethically justified according to consequentialist principles. In doing so, it contributes to a deeper understanding of how technological innovation intersects with moral responsibility, demonstrating that ethical governance is not merely a regulatory requirement but a moral imperative. The chapter argues that AI recruitment is ethically defensible only if its outcomes demonstrably promote social welfare, uphold equality, and respect the inherent dignity of all individuals involved in the hiring process.

5.1 Benefits of Using AI in Recruitment

Artificial Intelligence (AI) recruitment systems have significantly transformed human resource management by automating, optimising, and personalising numerous stages of the hiring process, ranging from advertising vacancies and screening résumés to conducting interviews and ranking candidates. Driver (2015, p. 97) argues that by handling large volumes of applications and standardising evaluation criteria, AI can address long-standing inefficiencies and bottlenecks inherent in traditional recruitment processes. From a consequentialist perspective, the ethical justification of these innovations rests on their outcomes: AI in

recruitment is morally desirable to the extent that it produces the greatest balance of good over harm for all affected parties. The adoption of AI recruitment systems can yield multiple positive outcomes. First, increased efficiency enables organisations to fill positions more quickly and accurately, thereby reducing delays that can negatively impact both employers and job seekers. The automation of repetitive tasks frees HR professionals to focus on strategic decision-making and employee engagement, thereby enhancing overall organisational effectiveness. Second, the reduction of human bias is a significant ethical benefit. By standardising candidate evaluation based on data-driven criteria, AI systems can minimise subjective judgments, personal prejudices, and unconscious discrimination, thereby promoting fairness and equal opportunity in the hiring process.

Third, cost savings represent another consequentially positive outcome. As Mejía (2019, pp. 312–314) cautions, AI reduces the financial and time burdens associated with manual screening, lengthy interview processes, and administrative overheads, allowing organisations to allocate resources more efficiently and potentially expand their workforce or invest in employee development initiatives. Moreover, AI recruitment can enhance access to employment by identifying suitable candidates from a larger and more diverse talent pool, including those who may be overlooked in conventional recruitment practices. This broader reach contributes to social welfare by improving employment opportunities and enabling economic participation for a wider range of individuals. Moreover, when implemented responsibly, these systems can simultaneously advance organisational objectives and broader socio-economic goals, thereby producing outcomes that consequentialist ethics would regard as morally commendable. The moral value of AI in recruitment lies not merely in its technological novelty but in its capacity to generate tangible benefits that advance both organisational performance and societal progress. When efficiency gains, fairness enhancements, and cost reductions are realised in a balanced and ethically informed manner, AI recruitment systems exemplify a consequentialist approach to technology adoption, prioritising actions that maximise positive outcomes while minimising harm for all stakeholders involved.

5.1.1 Increased Efficiency in Recruitment

Efficiency is often cited as the foremost advantage of AI in recruitment, as algorithms can process vast volumes of applications, shortlist qualified candidates, and match competencies to job specifications far more quickly than traditional manual methods. Chamorro Premuzic et al. (2020, p. 78) note that AI tools can reduce time to hire by up to 60 percent, enabling

organisations to fill vacancies faster and respond to market demands more effectively. In practical terms, recent industry reports indicate that organisations using AI-powered recruitment tools have achieved reductions in time to hire of up to 75 percent, particularly in high-volume hiring contexts. For example, HireSales (2025) platform reported that screening time for candidates dropped from 90 days to just 14 days after implementing AI-driven interview automation, representing a sixfold improvement. Similarly, a survey of Talent Genie (2025) found that 89 percent identified “saving time” as a primary benefit of AI in recruitment, highlighting its role in maintaining organisational productivity. Specific platforms, such as Workday (2024), have documented a 50 percent reduction in time-to-hire by tracking bottlenecks, including candidate drop-off and interview scheduling delays, and applying AI optimizations at each stage. While Talent Genie (2025) argues that automated screening tools report saving between 75–85 percent of the time normally spent on initial candidate evaluation, allowing human resource personnel to focus on strategic and people-centred activities such as candidate engagement and workforce planning.

From a consequentialist perspective, these efficiency gains carry moral significance only when they contribute to greater welfare for all stakeholders. In South Africa, where unemployment is high and administrative bottlenecks often slow down recruitment, reducing time-to-hire can produce substantial social benefits. Faster recruitment reduces the duration that job seekers remain unemployed, improving household income stability and overall social well-being. At the organisational level, increased efficiency enhances service delivery, particularly in public institutions or critical sectors, by ensuring that key roles are filled promptly, reducing operational gaps, and supporting smoother institutional functioning. However, efficiency alone is not sufficient to justify AI-driven recruitment from an ethical standpoint. As Mejía (2019, pp. 312–314) cautions, time savings and operational gains must not override fairness or human dignity.

AI systems that expedite recruitment but exclude marginalised groups, overburden candidates with algorithmic tests, or treat applicants as mere data points may increase efficiency yet produce ethically problematic outcomes. Likewise, a Workday 2024 study highlighted that candidates value transparent communication, fair assessment of soft skills, and respect for individual circumstances, demonstrating the ongoing importance of human judgment in the recruitment process. When implemented responsibly, AI-enabled efficiency aligns with consequentialist ethics by reducing individual suffering, enhancing organisational productivity, and freeing human resource professionals to engage in higher-value, human-centred work. In

the South African context, where unemployment and recruitment delays are pressing challenges, these outcomes contribute directly to social welfare, income security, and institutional effectiveness. Nevertheless, the ethical value of efficiency depends on how gains are achieved and who benefits. Efficiency that disproportionately advantages organisations while shifting burdens onto candidates or society would be morally weaker. In summary, AI-driven efficiency in recruitment represents a consequential positive development, provided it is paired with safeguards for fairness, transparency, and human dignity, ensuring that operational gains translate into broader social and economic benefits.

5.1.2 Reduced Human Bias

A significant ethical advantage of Artificial Intelligence (AI) in recruitment lies in its potential to mitigate human bias, a persistent issue in traditional hiring practices. Bogen and Rieke (2020, p. 15) observe that conventional recruitment processes are frequently shaped by unconscious biases related to race, gender, age, educational background, and geographic origin, which often distort merit-based evaluation and perpetuate exclusion. AI systems, when responsibly designed and trained on diverse and representative datasets, can evaluate candidates based on objective criteria such as skills, experience, qualifications, and performance indicators, rather than subjective impressions or personal preferences, as mentioned by Raghavan et al (2020, p. 470). This shift promotes fairness, transparency, and equal opportunity, generating outcomes that consequentialist ethics regard as morally desirable because they reduce discriminatory harm and enhance inclusion (Segall, 2021, p. 133). In the South African context, the Department of Employment and Labour (2023) states that where employment equity and redress remain central policy objectives, the ability of AI to counteract personal prejudice aligns closely with both constitutional mandates and moral imperatives. Mhlongo and Mabuza (2022, pp. 1–10) argue that AI recruitment tools can assist organisations in identifying qualified candidates from historically disadvantaged groups, rural areas, and underrepresented communities, thereby supporting the collective societal good by widening participation in the labour market and promoting socio-economic inclusion.

Several locally available platforms illustrate how bias reduction can be operationalised. For example, Talent Genie (2025) uses machine learning to match candidate profiles to job specifications, drawing on cross-cultural data and minimising reliance on subjective hiring cues. HireSales (2025) incorporates diversity metrics and long-term potential into its matching algorithm, enabling organisations to cast a wider net rather than reproducing narrow hiring models. Workday (2024) offers built-in bias-audit and transparency controls, signalling a move

toward governance and oversight in algorithmic decision-making. These tools demonstrate that AI recruitment in South Africa is not solely about speed or efficiency; it also holds the potential to advance structural fairness and support transformative hiring practices. However, the ethical value of AI in mitigating bias is not automatic. Hunkenschroer and Luetge (2022, p. 980) caution that such benefits are contingent on proactive governance, ongoing evaluation, and accountability mechanisms. Consequentialist ethics emphasises that positive outcomes, such as fairer hiring and broader inclusion, require empirical verification. Owolabi (2023), in a recent African study, highlights that AI systems trained on non-diverse datasets, often sourced from Western contexts, may inadvertently perpetuate discrimination and skew outcomes against underrepresented groups.

In South Africa, Cliffe Dekker Hofmeyr (2025) notes that while AI-driven recruitment raises hope for fairness, the absence of specific regulations governing algorithmic bias means that employers must remain vigilant. While Binns (2020, p. 520) and Raghavan et al. (2020, p. 479) note that without regular auditing, inclusive training data, and transparent governance, AI may replicate or even amplify existing biases, transforming a potentially beneficial tool into one that reinforces exclusion and inequality. Consequently, the ethical worth of AI in reducing human bias depends on how these systems are built, deployed, and monitored, and who benefits. When responsibly implemented, AI recruitment systems demonstrate a clear alignment with consequentialist principles: they maximise benefits such as justice, inclusion, and trust, while minimising harm caused by discriminatory practices. In the South African context, such systems not only enhance organisational fairness and efficiency but also contribute to broader social welfare and the ethical imperative of equitable access to employment opportunities.

5.1.3 Cost Savings in Recruitment

Artificial Intelligence (AI) recruitment systems offer substantial financial and operational benefits by reducing the administrative and logistical costs associated with traditional hiring. Tambe et al. (2019, p. 34) argue that technologies such as automated résumé screening, chatbots, and predictive analytics streamline tasks that would otherwise require extensive human labour and time. These efficiencies enable organisations to allocate resources more strategically, potentially lowering hiring costs while expanding recruitment reach to include a larger and more diverse pool of candidates. From a consequentialist perspective, cost efficiency holds ethical significance because it facilitates the redistribution of resources toward socially beneficial ends. As Driver (2021, p. 15) argues, savings generated through AI can be reinvested

in employee training, development programmes, or fairer remuneration structures, thereby enhancing both organisational performance and societal welfare. In South Africa's resource-constrained economy, reduced recruitment expenditure can improve organisational sustainability and competitiveness, while also supporting job creation and inclusive growth. The broader social consequence of AI-enabled cost savings is an increase in economic welfare through the more effective use of limited capital. Organisations may be better positioned to invest in transformation initiatives, expand access to employment, and contribute to national development goals. In this way, AI recruitment aligns with consequentialist ethics by producing outcomes that maximise collective good.

However, consequentialism also cautions that cost savings must not come at the expense of ethical integrity. If cost-cutting measures result in dehumanised candidate interactions, diminished user experience, or the exclusion of individuals lacking digital access, the negative consequences may outweigh the benefits. In South Africa, where digital divides persist across rural and urban populations, such risks are particularly salient. As Singer (2015, p. 40) notes, AI-driven cost efficiency carries moral worth only when it contributes to both organisational effectiveness and the collective good. This means improving recruitment outcomes while preserving fairness, transparency, and human dignity. Ethical cost savings are not merely about reducing expenditure; they are about enhancing value in ways that benefit all stakeholders, especially those historically marginalised in the labour market. AI recruitment's ability to generate cost savings is ethically commendable when those savings are responsibly reinvested and do not compromise the quality, inclusivity, or humanity of the hiring process. When implemented with care and accountability, these systems can support both economic efficiency and social justice, fulfilling the consequentialist imperative to maximise good while minimising harm.

5.2 Challenges of using AI in Recruitment

While Artificial Intelligence (AI) holds considerable promise for enhancing fairness, efficiency, and objectivity in recruitment, a consequentialist ethical analysis demands that these technologies be evaluated not by their intentions but by their outcomes. As Driver (2015, pp. 97–98) asserts, the moral legitimacy of an action depends on whether it produces more good than harm for all affected stakeholders. Accordingly, the ethical assessment of AI recruitment must critically examine its potential adverse effects, particularly those that diminish social welfare, perpetuate inequality, or erode human dignity. In the South African context, these concerns are especially acute. The country's labour market continues to bear the imprint of

racial exclusion, systemic inequality, and socio-economic stratification, despite ongoing efforts toward employment equity and transformation. If AI recruitment tools inadvertently reproduce or reinforce the very patterns of bias and exclusion that the Employment Equity Act 55 of 1998 seeks to dismantle, they cannot be ethically justified under consequentialist principles.

Moreover, the digital divide, disparities in access to technology, and the underrepresentation of African languages and cultural contexts in training data raise additional concerns about equity and inclusivity. In such a setting, the deployment of AI in recruitment must be approached with ethical vigilance, ensuring that its implementation does not exacerbate existing injustices. The following subsections outline four key ethical and practical challenges associated with AI recruitment: Algorithmic Bias and Discrimination explore how biased training data and opaque algorithms can lead to discriminatory outcomes, undermining fairness and reinforcing structural inequalities. Lack of Transparency and Accountability examines the ethical risks posed by “black box” decision-making, where candidates and employers alike may struggle to understand or challenge AI-driven outcomes. Dehumanisation of the Hiring Process that considers the moral implications of reducing human interaction in recruitment, including the erosion of empathy, dignity, and relational trust. Lastly, over-reliance on Technology, which analyses the risks of excessive dependence on AI tools, including diminished human oversight, loss of contextual judgment, and vulnerability to system failures. Each of these challenges is examined through a consequentialist lens, weighing its impact on fairness, inclusion, and social welfare. The goal is not to reject AI recruitment outright, but to ensure that its deployment is ethically defensible, contextually appropriate, and aligned with South Africa’s broader goals of justice, redress, and human dignity.

5.2.1 Algorithmic Bias and Discrimination in AI Recruitment

Algorithmic bias represents one of the most profound ethical challenges in AI-driven recruitment. As Bolukbasi, Chang, Zou, Saligrama, and Kalai (2019, pp.4349–4357) explain, such bias arises when machine-learning systems trained on historical hiring data reproduce or amplify patterns of discrimination embedded in those datasets. For instance, if past recruitment practices favoured certain racial, gender, or socio-economic groups, AI models may inadvertently “learn” to replicate these inequities, thereby perpetuating disadvantage for historically marginalised candidates. This undermines not only fairness and equity but also the constitutional right to equality enshrined in Section 9 of the Constitution of the Republic of South Africa (1996). From Segall's (2021, p. 133) consequentialist perspective, algorithmic bias constitutes a moral failure because it produces harmful outcomes such as exclusion,

reduced opportunity, and the reinforcement of structural inequality that outweigh potential benefits like efficiency or objectivity. Mejía (2019, pp. 312–314) further argues that an ethical system must be evaluated by how it distributes benefits and burdens across society. When AI algorithms systematically exclude candidates from historically disadvantaged universities, rural communities, or those with informal or non-traditional work experience, they contribute to social harm rather than collective welfare.

In South Africa, where employment is a critical pathway to socio-economic empowerment, algorithmic discrimination is not merely a technical flaw; it is a moral and developmental crisis. It directly contravenes the objectives of the Employment Equity Act 55 of 1998 and the Broad-Based Black Economic Empowerment (B-BBEE) framework, as mentioned in the Department of Employment and Labour (2023) and South African Labour Guide (2022, p. 8), both of which aim to redress historical exclusion and promote equitable workforce participation. As Hunkenschroer and Luetge (2022, p. 980) and Raghavan et al. (2020, pp. 469–481) argue that consequentialist ethics, therefore, demand that AI recruitment systems undergo continuous auditing, bias testing, and retraining to ensure that their outcomes support social justice and inclusion rather than exclusion. This includes not only technical adjustments but also institutional accountability and ethical oversight. For example, Talent Genie (2025), a local AI recruitment platform, integrates cross-cultural and representative data to mitigate bias in candidate shortlisting.

Whereas Workday (2024) includes bias-auditing tools that allow recruiters to identify and correct inequitable algorithmic patterns, promoting transparency and fairness in decision-making. These measures demonstrate practical approaches to aligning AI outcomes with ethical and legal imperatives, ensuring that efficiency and objectivity do not come at the cost of fairness, inclusion, or societal welfare. Algorithmic bias in AI recruitment is both a technical and ethical challenge. From a consequentialist standpoint, its moral legitimacy depends entirely on whether its outcomes reduce harm, promote equity, and advance social welfare. In South Africa, where historical and structural inequalities remain pervasive, addressing algorithmic bias is essential to ensuring that AI recruitment contributes positively to national development goals, workforce transformation, and the equitable distribution of opportunity.

5.2.2 Lack of Transparency and Accountability in AI Recruitment

Transparency and accountability are foundational to ethical decision-making, particularly when automated systems influence human lives and livelihoods. In the context of AI-driven

recruitment, these principles become even more critical. Raghavan, Barocas, Kleinberg, and Levy (2020, p. 470) observe that many AI recruitment systems operate as “black boxes”, meaning their internal logic, weighting of criteria, and decision-making processes are inaccessible to employers, candidates, and regulators. Binns (2020, p. 520) argues that this opacity prevents stakeholders from understanding, challenging, or contesting hiring decisions, thereby undermining procedural fairness. From a consequentialist perspective, such a lack of transparency produces negative moral consequences: it erodes trust, increases the risk of injustice, and diminishes institutional legitimacy. The South African Department of Employment and Labour (2023) notes that candidates who are unaware of why they were rejected may experience emotional distress, reduced self-efficacy, and alienation, while employers may struggle to ensure compliance with legal frameworks such as the Protection of Personal Information Act 4 of 2013 (POPIA) and the Promotion of Administrative Justice Act 3 of 2000. Mejía (2019, pp. 312–314) advocates for an institution-sensitive consequentialist approach, which evaluates ethical legitimacy by considering the systemic effects of institutional practices. In this light, unaccountable AI recruitment systems produce collective harm by weakening oversight mechanisms and eroding the procedural integrity of hiring processes.

In South Africa, where public trust in institutions remains fragile and where employers are expected to demonstrate commitment to transformation imperatives, opaque AI systems may further reduce confidence in both public and private recruitment. As reported by News24 (2024), the use of AI without transparency risks alienating job seekers and undermining perceptions of fairness. Cliffe Dekker Hofmeyr (2025) further warns that employers using AI without clear documentation or explainability may inadvertently expose themselves to claims of unfair discrimination or non-compliance with employment equity objectives. As Hunkenschroer and Luetge (2022, p. 980) and Segall (2021, p. 133) note, consequentialist ethics demand that AI recruitment systems be transparent about their data sources, algorithmic logic, and decision criteria, while ensuring that human recruiters remain accountable for final hiring outcomes. Practical measures to uphold these standards include providing candidates with clear explanations for rejections, documenting decision-making workflows, performing regular algorithmic audits, and retaining human oversight to mitigate unintended consequences.

In South Africa, platforms such as Accely Recruitment AI, Talent Genie (2025), and Workday (2024) have incorporated bias-auditing features and reporting dashboards that enable HR

personnel to monitor and adjust algorithmic decisions. These tools demonstrate responsible implementation aligned with both legal compliance and ethical best practices. By promoting transparency and accountability, such measures generate morally desirable outcomes, including fairness, trust, and institutional credibility, all of which are key indicators of ethical legitimacy under a consequentialist framework. AI recruitment systems must not only function efficiently but also operate in ways that are understandable, contestable, and justifiable, thereby reinforcing public confidence and advancing social welfare.

5.2.3 Dehumanization of the Recruitment Process

A growing ethical concern in AI-driven recruitment is the dehumanisation of candidates, the reduction of individuals to data points, algorithmic profiles, and predictive scores, rather than recognising them as autonomous human beings with unique potential. Newlands (2021, p. 92) highlights that automated assessments, video interviews analysed by facial-recognition software, and personality predictions based on linguistic patterns can strip away empathy, dialogue, and context, resulting in an impersonal and mechanised hiring experience. From a consequentialist standpoint, moral worth must be measured by the impact on well-being and dignity. As Driver (2021, p. 15) argues, when AI recruitment erodes human connection or causes psychological harm, making candidates feel objectified, powerless, or alienated, it produces outcomes inconsistent with the principle of welfare maximisation. In South Africa, where human dignity is constitutionally enshrined in the Constitution of the Republic of South Africa (1996, s. 10), this challenge carries profound moral and legal weight.

The Department of Employment and Labour (2023) notes that for many job seekers, particularly those in under-resourced or rural communities with limited digital access, interacting with impersonal AI systems may exacerbate feelings of exclusion, confusion, and disadvantage. Segall (2021, p. 133) asserts that ethical systems must account for structural injustice; dehumanising recruitment processes are morally impermissible because they reproduce social hierarchies under a veneer of technological neutrality. Furthermore, Mhlongo and Mabuza (2022, pp.1–10) report that South African candidates often experience stress, distrust, and emotional discomfort when subjected to AI-driven evaluations without human explanation or interaction. These experiences undermine not only candidate confidence but also public trust in recruitment systems and institutional fairness. Consequentialism, therefore, demands that efficiency and automation must never come at the expense of human welfare. Maintaining empathy, providing candidate feedback, and ensuring human oversight are essential to safeguarding dignity and well-being. Segall (2021, p. 133) and Chamorro-Premuzic

et al. (2020, pp.69–73) argue that by integrating human judgment alongside AI systems, organisations can ensure that technology serves human needs, rather than replacing moral and relational considerations. In this way, AI recruitment can align operational efficiency with ethical responsibility, ensuring that candidate dignity and welfare remain central to the hiring process. When human oversight complements algorithmic accuracy, recruitment becomes not only more effective but also more just, inclusive, and trustworthy, outcomes that consequentialist ethics would regard as morally commendable.

5.2.4 Over-Reliance on Technology

The increasing dependence on Artificial Intelligence (AI) for decision-making in recruitment raises the ethical concern of over-reliance on technology, a condition in which human evaluators defer excessively to algorithmic outputs without critical scrutiny. Hunkenschroer and Luetge (2022, p.985) warn that such reliance can erode human judgment and moral responsibility, particularly when recruiters assume that algorithmic recommendations are inherently objective or superior. Raghavan, Barocas, Kleinberg, and Levy (2020, pp.469–481) describe this phenomenon as “automation bias”, which fosters moral passivity and diminishes the capacity for ethical reflection. From a consequentialist perspective, over-reliance is ethically problematic because it may produce harmful outcomes that outweigh the benefits of efficiency. As Driver (2015, pp. 97–98) argues, when organisations outsource ethical judgment to machines, they risk enabling unchecked discrimination, exclusion, and systemic bias, thereby negatively impacting candidates’ well-being, and undermining societal welfare. In the South African labour context, over-reliance on AI may conflict with legal and regulatory frameworks designed to protect fairness and accountability. The Protection of Personal Information Act 4 of 2013 (POPIA, s. 71) prohibits decisions based solely on automated processing unless safeguards such as human review are implemented. Similarly, the Employment Equity Act 55 of 1998 (EEA) and the Promotion of Administrative Justice Act 3 of 2000 (PAJA) emphasise procedural fairness, reasonableness, and accountability requirements that cannot be guaranteed without human oversight.

Over-reliance on technology can therefore reinforce exclusion, exacerbate inequality, and allow systemic bias to persist, especially when algorithms are treated as infallible. Consequentialist ethics supports a hybrid decision-making model in which AI augments, rather than replaces, human judgment. Segall (2021, p. 133) argues that this approach ensures that algorithmic efficiency and predictive accuracy are balanced with moral responsibility, enabling employers to critically evaluate outputs, provide candidate feedback, and intervene where

necessary to prevent harm. In practice, South African organisations have begun incorporating such safeguards. For example, Talent Genie (2025) and Accely Recruitment AI offer interactive dashboards that allow HR personnel to monitor and adjust algorithmic recommendations. These platforms ensure that human accountability remains central, enabling recruiters to uphold ethical standards while benefiting from technological support. By maintaining this balance, employers preserve moral agency and ensure that recruitment outcomes remain fair, transparent, and socially beneficial, consistent with the principles of consequentialist ethics. AI should be a tool that enhances human judgment, not one that replaces it. Ethical recruitment requires that technology serve human values, ensuring that efficiency does not come at the cost of justice, dignity, or inclusion.

5.3 The Theory of Consequentialism Framework

Consequentialist ethics provides the evaluative framework for assessing the moral legitimacy of AI recruitment systems. According to this framework, the moral value of an action depends entirely on whether it produces more good than harm for all affected parties (Driver, 2015). Applied to AI recruitment, this means the systems are ethically justified only when their outcomes generate greater welfare, justice, and human dignity while minimizing harm and discrimination.

5.3.1 Table 1: Ethical Assessment of Benefits of Using AI in Recruitment

Benefit	Moral Significance	Welfare Impact	Conditions
Efficiency Gains	Reduces individual suffering	Faster employment, improved productivity	Must pair with fairness safeguards
Bias Reduction	Advances justice and inclusion	Expands participation, promotes equity	Requires empirical verification
Cost Savings	Enables resource reallocation	Supports sustainability and job creation	Must not compromise dignity

Own source (2025)

The table titled “Ethical Assessment of Benefits” evaluates the moral and welfare implications of adopting Artificial Intelligence (AI) in recruitment. It aligns with the objective “to assess the ethical consequences of AI adoption in recruitment” by systematically analysing how AI-

driven practices can yield moral and social outcomes while identifying the necessary conditions for ethical implementation.

The first benefit, Efficiency Gains, refers to AI's ability to streamline recruitment by automating repetitive tasks, reducing hiring time, and improving matching accuracy. Floridi and Cowls (2019) observe that, ethically, this reduces individual suffering associated with prolonged job searches and organizational inefficiency. The moral significance lies in promoting human welfare through faster employment and productivity. However, the text emphasizes that efficiency must be balanced with fairness safeguards to ensure that speed and automation do not perpetuate or amplify discrimination.

The second benefit, Bias Reduction, addresses AI's potential to advance justice and inclusion. When professionally designed and tested (Binns, 2018), observing AI can help eliminate unconscious human biases in hiring, thus expanding participation and promoting equity. However, this moral gain requires empirical verification, as AI systems trained on biased data may unintentionally replicate systemic inequalities. Continuous auditing and transparency are thus essential ethical conditions.

The third benefit, Cost Savings, highlights AI's ability to reallocate resources by automating administrative tasks and reducing recruitment expenses. This can support sustainability and create new employment opportunities in tech and data sectors. Yet, IEEE (2021) emphasizes the ethical condition, stresses that economic efficiency must not compromise human dignity; workers and candidates should still be treated with respect and given meaningful engagement throughout the hiring process.

Overall, the table highlights that while AI adoption in recruitment offers substantial moral and welfare benefits, these outcomes depend on rigorous ethical governance that ensures fairness, transparency, and respect for human dignity.

5.3.2 Table 2: Ethical Assessment of Challenges of Using AI in Recruitment

Challenge	Harm	Justice Impact	Ethical Status
Algorithmic Bias	Reinforces exclusion and inequality	Contradicts EEA objectives	Negative
Lack of Transparency	Undermines trust and legitimacy	Erodes procedural fairness	Negative
Dehumanization	Causes psychological harm	Reproduces social hierarchies	Negative
Over-Reliance	Diminishes moral responsibility	Violates legal frameworks	Negative

Own source (2025)

The table titled “Ethical Assessment of Challenges” critically examines the negative ethical consequences associated with AI adoption in recruitment, linking to the objective “to assess the ethical consequences of AI adoption in recruitment.” It identifies key challenges algorithmic bias, lack of transparency, dehumanisation, and over-reliance, which pose significant moral and justice-related concerns, each marked with a negative ethical status due to their potential to undermine fairness, dignity, and accountability.

The first challenge, Algorithmic Bias, occurs when AI systems replicate or amplify historical and social inequalities embedded in training data. This bias reinforces exclusion and inequality, leading to discriminatory hiring outcomes that contradict the Equal Employment and Anti-Discrimination (EEA) objectives (Barocas and Selbst, 2016). Ethically, such outcomes violate principles of justice and fairness, making algorithmic bias a serious negative concern unless mitigated through data auditing and diverse model training.

Lack of Transparency represents another ethical challenge, as Diakopoulos (2016) notes opaque algorithms prevent candidates and employers from understanding how recruitment decisions are made. This undermines trust and legitimacy and erodes procedural fairness, essential components of ethical governance. Transparency and explainability are thus necessary to maintain moral accountability in AI-driven hiring.

The third challenge, Dehumanization, refers to the reduction of candidates to data points, causing psychological harm by stripping away individuality and empathy in evaluation

processes (Turkle, 2017). This practice perpetuates social hierarchies and undermines the moral fabric of recruitment by prioritising efficiency over human dignity.

Finally, Over-Reliance on AI can lead to a diminished sense of moral responsibility, as decision-makers defer too heavily to algorithmic outcomes. This not only violates the legal and ethical frameworks that require human oversight but also raises concerns about accountability when errors occur (European Commission, 2021). Overall, the table highlights that these challenges undermine justice and well-being, reinforcing the need for ethical safeguards, human oversight, and transparent design to ensure AI recruitment aligns with moral and legal standards.

5.8 Conclusion

This chapter has critically examined the ethical implications of Artificial Intelligence (AI) in recruitment and selection through the lens of consequentialism, a moral theory that evaluates actions based on their outcomes rather than their intentions. Guided by the principle that morally right actions are those that produce the greatest balance of good over harm, the analysis has demonstrated that AI recruitment presents both significant ethical benefits and serious moral risks, particularly within the South African socio-economic context. The discussion highlighted that AI systems can generate positive moral outcomes by enhancing efficiency, minimising human bias, and reducing recruitment costs. From a consequentialist standpoint, these benefits carry ethical weight because they contribute to social welfare, accelerate employment processes, promote equity, and enable a more effective allocation of scarce organisational resources. In a society marked by persistent unemployment and inequality, such improvements can enhance collective well-being by supporting inclusive economic participation. When responsibly implemented, AI thus emerges not merely as a business tool, but as an instrument of ethical and developmental value.

Although the central focus of this dissertation is ethical, the analysis necessarily engages with legal, regulatory, and empirical paradigms. This is because ethical evaluation cannot be divorced from the institutional and empirical realities in which AI recruitment systems operate. Legal frameworks shape the boundaries of permissible practice, while empirical data illustrate the lived consequences of algorithmic decision-making. By situating ethical inquiry alongside these dimensions, the study demonstrates how normative ideals of fairness, equality, and dignity intersect with regulatory compliance and real-world outcomes. This integrative approach ensures that the ethical critique is not abstracted from the South African labour

context but instead grounded in the socio-economic and institutional structures that influence recruitment practices.

However, consequentialist reasoning also reveals that the ethical legitimacy of AI recruitment is conditional upon its real-world consequences. Algorithmic bias, opacity, dehumanisation, and excessive reliance on automation can all produce harmful outcomes that outweigh efficiency gains. These harms manifest as discrimination, erosion of human dignity, and diminished accountability outcomes that are incompatible with both consequentialist morality and South Africa's constitutional commitment to equality and fairness. The analysis, therefore, underscores a critical insight: technological innovation is not ethically neutral. AI systems must be deliberately designed, governed, and monitored to promote justice, inclusion, and transparency, rather than entrench systemic inequality. In practical terms, consequentialism demands the continuous evaluation of AI's net social outcomes. Developers, employers, and policymakers must ensure that algorithmic design, data governance, and implementation frameworks actively prevent bias, uphold human dignity, and protect the rights of candidates.

This chapter has demonstrated that consequentialism offers a rigorous and contextually relevant ethical framework for evaluating AI recruitment within South Africa's institutional and historical landscape. Nevertheless, through weighing both the benefits and risks of algorithmic hiring across dimensions of efficiency, fairness, transparency, dignity, and societal impact, the study has revealed both the promise and the peril of technological automation in the employment context. Consistent with Mejía's institution-sensitive consequentialism and Segall's justice-oriented approach, the findings suggest that AI recruitment can be morally justified only when its real-world consequences demonstrably enhance welfare, equality, and human dignity. This outcome-based evaluation provides a clear foundation for Chapter Six, which will synthesise the research insights into ethical conclusions and actionable recommendations. The next chapter thus moves from analysis to application, translating consequentialist reasoning into policy and governance strategies that promote responsible, inclusive, and equitable AI use in recruitment.

CHAPTER SIX: SUMMARY, CONCLUSION, AND RECOMMENDATIONS

6.0 Introduction

This concluding chapter draws together the key insights generated throughout the study and situates them within the wider ethical, technological, and socio-economic debates surrounding Artificial Intelligence (AI) in recruitment and selection. It offers an integrated synthesis of how the research has met its objectives and responded to the central research question: What are the ethical effects of using AI in job recruitment and selection within the South African context when examined through the lens of consequentialist theory? In doing so, the chapter provides a reflective overview of the main findings, clarifies their significance for ethical theory and professional practice, and highlights their implications for future policy and scholarship. The discussion begins by reaffirming the study's grounding in consequentialist ethics, a normative theory that measures moral worth through the balance of good and adverse consequences produced by an action. As established by Driver (2015, pp. 97–98) and extended by de Lazari-Radek and Singer (2017, pp. 62–65), consequentialism evaluates technological innovations not by their intentions or internal logic but by the extent to which they improve human welfare and reduce harm. This theoretical lens has guided the study's interpretation of AI recruitment outcomes, allowing for a systematic assessment of both efficiency gains and ethical risks.

Within South Africa's complex socio-economic landscape marked by historical inequality, structural unemployment, and aspirations for inclusive growth, the ethical evaluation of AI recruitment acquires urgency. The diffusion of algorithmic hiring technologies offers potential for transparency, speed, and fairness, yet these systems also risk amplifying bias, marginalisation, and loss of dignity if implemented without ethical safeguards. Consequentialism provides an effective analytical framework for navigating these tensions because it focuses on outcomes that matter to people's lives: equitable access to work, institutional trust, and collective well-being. This chapter, therefore, performs several interrelated functions. It synthesises the theoretical and empirical findings of the research, demonstrating how each has contributed to answering the guiding questions. It then extrapolates from those findings to propose normative conclusions about the moral legitimacy of AI recruitment practices in South Africa. Finally, it translates the study's insights into practical and policy-oriented recommendations aimed at promoting ethical AI governance, institutional accountability, and social justice.

By integrating the analytical results with the broader ethical discourse, the conclusion moves beyond summary to critical reflection. It recognises that technology is never morally neutral;

its value depends on the consequences it produces for human beings and societies. In this sense, the chapter affirms that AI's role in recruitment must be judged not solely by its technical sophistication or efficiency but by its contribution to human flourishing, dignity, and fairness. Accordingly, the sections that follow 6.1 Summary, 6.2 Conclusion, and 6.3 Recommendations progress from synthesising the research findings to evaluating their ethical implications and proposing pathways for the responsible use of AI in South Africa's evolving labour landscape.

6.1 Summary of the Study

This study set out to critically examine the ethical effects of using Artificial Intelligence (AI) in job recruitment and selection within the South African context, employing the analytical lens of consequentialist ethics. The research aimed to evaluate whether AI-driven recruitment systems promote social welfare, fairness, and human dignity, or whether they carry the potential to reinforce bias, exclusion, and systemic inequality. Guided by the central research question: To what extent can the use of AI in recruitment and selection be considered ethically justified when evaluated according to the principles of consequentialism? the study explored both the positive and negative consequences of AI implementation in hiring practices. Through a detailed literature review and methodological analysis, the study identified key benefits of AI in recruitment, including increased efficiency, reduced human bias, and cost savings, all of which can enhance organisational performance and contribute to broader social welfare when deployed responsibly (Chamorro-Premuzic, Winsborough, Sherman, and Hogan, 2020; Bogen and Rieke, 2020; Driver, 2015, pp. 97–98).

Conversely, the study highlighted several ethical and practical challenges that may undermine these benefits, such as algorithmic bias, lack of transparency and accountability, dehumanisation of candidates, and over-reliance on technology (Bolukbasi et al., 2019; Hunkenschroer and Luetge, 2022; Newlands, 2021; Raghavan et al., 2020). These challenges are particularly significant in the South African labour market, where historical inequalities, constitutional mandates, and regulatory frameworks like the Employment Equity Act (EEA) and the Protection of Personal Information Act (POPIA) necessitate equitable and accountable hiring practices (Department of Employment and Labour, 2023; South African Labour Guide, 2022). Overall, the study demonstrates that the ethical justification of AI in recruitment depends on the actual consequences of its deployment. A consequentialist perspective requires that AI systems not only improve operational efficiency but also actively enhance fairness, inclusion, and human dignity, while minimising harm to candidates and society. The findings provide a foundation for formulating practical recommendations that balance technological

innovation with ethical and legal responsibilities, guiding organisations in implementing AI recruitment systems that are morally defensible and socially beneficial.

Chapter One: Introduction

The opening chapter established the foundation of the study by providing a comprehensive overview of the research context, rationale, problem statement, aims, and objectives. It highlighted the growing global reliance on Artificial Intelligence (AI) in recruitment and selection processes, situating South Africa within this technological transformation and underscoring the nation's unique socio-economic and legal context. While AI offers the potential for enhanced efficiency, objectivity, and scalability in hiring practices, the chapter recognised its simultaneous risk of reproducing discriminatory patterns and undermining fairness if not ethically managed. The research problem was framed around the moral ambiguity of AI recruitment: specifically, whether automation enhances ethical hiring or compromises principles of justice and human dignity. The chapter outlined the research objectives to critically analyse the ethical consequences of AI-driven recruitment, evaluate these outcomes through the lens of consequentialist theory, and develop policy and practical recommendations to support the ethical implementation of AI in hiring. Furthermore, it highlighted the significance of the study in contributing to both academic discourse and practical governance. In the South African context, characterised by structural unemployment, economic inequality, and legislative imperatives such as the Employment Equity Act 55 of 1998 (EEA) and the Protection of Personal Information Act 4 of 2013 (POPIA), examining the moral implications of AI recruitment is particularly crucial for ensuring technological justice and promoting inclusive workforce participation. Finally, the chapter defined key concepts, including Artificial Intelligence, recruitment ethics, and consequentialism, to provide conceptual clarity for subsequent analysis. It also outlined the structure of the dissertation, signalling how each chapter would build on the preceding one to construct a comprehensive ethical evaluation of AI in recruitment and selection. The introduction thereby established a clear foundation for understanding both the opportunities and risks of AI adoption in South African recruitment practices, setting the stage for the literature review, methodology, and theoretical frameworks that follow.

Chapter Two: Literature Review

Chapter Two critically reviewed existing scholarship on Artificial Intelligence (AI), ethics, and employment to identify key debates, insights, and gaps in knowledge relevant to AI

recruitment. The review synthesised both international and South African literature, engaging with studies such as Ajunwa (2020, pp. 12–13), Bogen and Rieke (2020, p. 15), and Zou and Schiebinger (2021, p. 324), which examine algorithmic bias, discrimination, and ethical concerns in AI-driven recruitment systems. While global research has extensively explored the ethical implications of AI, the literature review revealed a relative paucity of studies contextualised within South Africa’s unique socio-political and labour market realities, where historical inequalities and legal mandates such as the Employment Equity Act (EEA) and Broad-Based Black Economic Empowerment (B-BBEE) frameworks shape the ethical stakes of automated hiring. The chapter also engaged with literature on ethical theories in technology governance, contrasting consequentialism, deontology, and virtue ethics, and justifying the adoption of consequentialism as the analytical lens for this study. Empirical research on AI’s effects on recruitment efficiency, fairness, and transparency was considered, drawing on the work of Chamorro-Premuzic, Winsborough, Sherman, and Hogan (2020, p. 78) and Tambe, Cappelli, and Yakubovich (2019, p. 34), highlighting both the opportunities and limitations of AI in different institutional and cultural contexts. Overall, Chapter Two demonstrated that while AI has the potential to transform recruitment practices, its ethical legitimacy depends primarily on the social consequences it produces. These insights directly informed the methodological design and theoretical framework of the study, providing a foundation for the consequentialist evaluation of AI recruitment in subsequent chapters.

Chapter Three: Research Method and Methodology

Chapter Three detailed the research design and methodological approach employed to explore the ethical consequences of AI recruitment in South Africa. The study adopted a qualitative, exploratory, desktop-based research design, which is particularly suitable for analysing the conceptual and normative dimensions of ethics (Lincoln and Guba, 1985, p. 290). This approach prioritises trustworthiness, contextual understanding, and critical interpretation over quantification, allowing for an in-depth evaluation of the moral implications of AI in recruitment. Data were collected from secondary sources, including peer-reviewed journal articles, policy and legislative documents, reports from organisations such as the World Economic Forum (2025) and UNESCO (2021), and case studies of AI recruitment platforms implemented in South African workplaces. Purposive sampling was employed to select literature directly relevant to AI ethics, recruitment fairness, and consequentialist evaluation, ensuring that the study focused on sources that could provide meaningful insights into the research question. The DECA method (Describe, Evaluate, Consult, Act) was used to structure

the data analysis, providing a systematic approach for assessing the outcomes of AI recruitment, evaluating their ethical significance, and consulting normative frameworks to draw actionable conclusions. The chapter also addressed strategies for ensuring trustworthiness and methodological rigour, including credibility, transferability, dependability, and confirmability, following Lincoln and Guba's (1985) criteria. Ethical considerations were carefully observed throughout, particularly in the objective interpretation of secondary data and the avoidance of personal or disciplinary bias. By combining a robust qualitative design with a structured analytic framework, the methodology provided a solid foundation for the theoretical and analytical discussions presented in subsequent chapters, enabling a consequentialist evaluation of AI recruitment systems grounded in both scholarly and practical evidence.

Chapter Four: Theoretical Framework: Consequentialism

Chapter Four presented the theoretical lens underpinning the study: consequentialism, an ethical framework that evaluates the morality of actions based on their outcomes rather than their intentions or adherence to prescriptive rules. Drawing on Driver (2015, pp. 97–98), the chapter defined consequentialism as the principle that morally right actions are those that produce the greatest good while minimising harm. The discussion traced the evolution of consequentialist thought from classical utilitarianism, as articulated by Bentham and Mill, to modern adaptations that emphasise justice, institutional sensitivity, and context-specific considerations (de Lazari-Radek and Singer, 2017, pp. 62–65; Mejía, 2019, pp. 312–314; Segall, 2021, p. 133; Arrhenius, Bykvist, and Campbell, 2020, pp. 115–120). These contemporary theorists expanded the framework to account not only for individual welfare but also for the broader societal and institutional consequences of moral actions, making consequentialism well-suited for analysing technology-driven decision-making. Importantly, the chapter contextualised these theoretical perspectives within South Africa's socio-economic and moral landscape, where structural inequality, unemployment, and historical exclusion demand that ethical evaluations consider systemic effects on fairness, inclusion, and collective well-being. By applying a consequentialist lens, the study framed AI recruitment systems as morally defensible only if their implementation generates positive outcomes, such as enhanced social welfare, equitable access to opportunities, and reduced discrimination, while minimising harm to individuals and communities. This theoretical foundation guided the subsequent chapters' analysis, providing the normative criteria for assessing the ethical legitimacy of AI-driven hiring practices in South Africa.

Chapter Five: Analysis and Discussion

Chapter Five applied the consequentialist ethical framework to critically evaluate the moral dimensions of AI-driven recruitment in the South African context. The chapter systematically analysed four core themes: operational efficiency, transparency and accountability, algorithmic bias and fairness, and human dignity and dehumanisation. Under operational efficiency, evidence from Chamorro-Premuzic et al. (2020, p. 78) and Tambe et al. (2019, p. 34) demonstrated that AI can enhance productivity, reduce time-to-hire, and streamline recruitment processes. From a consequentialist perspective, such efficiency is morally valuable only when it contributes to inclusive welfare by improving access to employment and reducing socio-economic disadvantage, rather than merely serving corporate interests. However, excessive reliance on automation, as noted by Newlands (2021, p. 92) and Segall (2021, p. 133), risks dehumanising recruitment processes and reinforcing structural inequality, highlighting the ethical tension between speed and social welfare. The theme of transparency and accountability examined the capacity of AI systems to enhance trust through explainable decision-making or to erode confidence when operations remain opaque (Goodman and Flaxman, 2017, p. 54; Keller and Aplin, 2024, pp. 9–11; Mejía, 2019, pp. 312–314). The analysis concluded that genuine transparency, combined with human oversight, maximises collective welfare by enabling stakeholders to understand and contest decisions, whereas superficial disclosure generates mistrust and diminishes institutional legitimacy. Regarding algorithmic bias and fairness, the study engaged with Ajunwa (2020, p. 1673), Bogen and Rieke (2020, p. 15), and Zou and Schiebinger (2021, p. 324) to demonstrate that biased datasets can perpetuate systemic injustice, particularly in societies like South Africa with entrenched structural inequalities. Consequentialist evaluation emphasises that fairness cannot be assumed from design intentions alone; it must be empirically verified through measurable outcomes that promote inclusion and equitable access to opportunity. Finally, the analysis of human dignity highlighted the risks associated with dehumanisation in AI recruitment, where candidates may be reduced to data points and predictive scores (Hunkenschroer and Luetge, 2022, p. 980). Consequentialist ethics underscores that moral legitimacy requires AI processes to preserve empathy, respect, and candidate autonomy, which are essential contributors to social welfare. In sum, Chapter Five concluded that AI recruitment holds substantial potential to improve efficiency, fairness, and inclusion, but its ethical validity depends on demonstrable outcomes that maximise welfare, minimise harm, and uphold human dignity. In the South African context, where historical inequalities and regulatory frameworks like the Employment Equity Act and POPIA shape the

labour landscape, the ethical deployment of AI necessitates careful design, continuous monitoring, and human oversight to ensure morally defensible and socially beneficial outcomes.

Chapter Six: Conclusion and Recommendations

The concluding chapter synthesises the study's findings, reaffirming that the ethical permissibility of AI in recruitment is contingent on its ability to produce more social benefit than harm. Evaluated through the consequentialist framework, the study emphasises that moral assessment must prioritise outcomes over intentions, focusing on whether AI systems advance fairness, inclusion, and collective well-being. The analysis throughout the dissertation demonstrated that while AI can improve efficiency, reduce human bias, and enhance organisational effectiveness, these benefits are ethically defensible only if they are accompanied by safeguards that prevent discrimination, dehumanisation, and over-reliance on technology. Based on the findings, the study proposes several key recommendations for ethical AI recruitment. First, organisations should implement robust governance frameworks that ensure human oversight, enabling recruiters to critically review algorithmic outputs and intervene when necessary. Second, transparency must be embedded in AI systems, including clear explanations of decision criteria, data sources, and algorithmic processes, to build trust and accountability among candidates and stakeholders. Third, algorithmic fairness should be actively monitored through continuous auditing, inclusive training datasets, and bias mitigation strategies to prevent the reproduction of structural inequities. Fourth, human dignity must be preserved by maintaining empathetic engagement, feedback mechanisms, and contextual evaluation alongside automated assessments. The chapter also identifies areas for further research, particularly regarding the long-term and intergenerational consequences of AI in employment, including its impact on workforce diversity, career progression, and societal equity. The study argues that technological innovation must be guided by ethical principles that prioritise human welfare, rather than allowing technology to dictate moral norms. Consequentialism provides a practical and justifiable framework for this purpose, compelling policymakers, human resource practitioners, and technology developers to ensure that AI-driven recruitment functions as a tool for empowerment, inclusion, and social good rather than exclusion. The dissertation underscores that the moral legitimacy of AI in recruitment is inseparable from its tangible effects on fairness, dignity, and collective well-being, offering actionable guidance for ethically responsible implementation in the South African context and beyond.

6.2 Conclusion

This study concludes that the ethical legitimacy of Artificial Intelligence (AI) in recruitment and selection depends fundamentally on the quality of its consequences rather than the sophistication of its algorithms or the intentions of its designers. From a consequentialist perspective, as articulated by Driver, moral rightness derives from the extent to which an action maximises good and minimises harm. Applied to AI recruitment, this means that technology must be evaluated by the welfare outcomes it produces for individuals, organisations, and society. The research has demonstrated that AI systems possess undeniable potential to enhance efficiency, objectivity, and fairness in recruitment. Evidence from Chamorro-Premuzic et al. and Tambe et al. shows that automation can significantly reduce administrative delays and align candidates more accurately with job requirements. These improvements generate positive utility when they expand opportunities, lower costs, and enable faster employment outcomes that resonate with consequentialist ideals of welfare maximisation. In contexts such as South Africa, where unemployment and skills mismatch remain critical issues, the responsible use of AI could materially contribute to social good by improving access to work and enhancing organisational effectiveness.

However, consequentialism demands a holistic assessment of all consequences, including negative and long-term ones. The findings revealed several ethically troubling outcomes that may outweigh the benefits if unmitigated. Ajunwa, Bogen, and Rieke have shown that biased training data and opaque algorithms often reproduce patterns of discrimination. In South Africa's post-apartheid labour market, such outcomes deepen structural inequalities rather than reduce them. Consequentialist reasoning, therefore, regards these as moral failures because they diminish collective welfare and perpetuate harm among vulnerable groups. The research also revealed that opacity in algorithmic decision-making undermines both fairness and trust. As Goodman and Flaxman and Raghavan et al. observe, candidates denied transparency or the ability to contest AI-based decisions experience procedural injustice and emotional distress. These harms carry measurable social costs, including diminished confidence in institutions and reduced motivation among job seekers. Under consequentialist ethics, such negative utility cannot be ignored; transparency and accountability thus become moral imperatives, not optional corporate virtues.

Furthermore, the study underscored the dehumanising potential of automated recruitment. Drawing on Hunkenschroer and Luetge it was found that when candidates are reduced to digital profiles and predictive scores, essential human qualities creativity, empathy, and resilience,

risk being overlooked. This mechanisation of human value violates the consequentialist obligation to preserve dignity and well-being as core components of utility. In this sense, the moral evaluation of AI must extend beyond quantitative outcomes to qualitative human experiences. Taken together, the findings affirm that the moral legitimacy of AI recruitment cannot be presumed from technological progress alone. Consequentialism, particularly in its institution-sensitive form as articulated by Mejía, insists that ethical evaluation must account for the institutional and historical contexts within which consequences unfold. In South Africa, with its legacy of inequality and ongoing struggle for social justice, AI recruitment becomes a test case for whether digital innovation can serve the public good rather than reinforce exclusion.

The study concludes that AI recruitment systems are ethically justified only when they demonstrably enhance human welfare across three interconnected dimensions: Individual well-being, by treating candidates fairly, transparently, and with respect; Organisational benefit, by promoting diversity, innovation, and effective matching; and Societal good, by reducing unemployment, building trust, and advancing equality. When these outcomes are realised, AI contributes positively to collective welfare and meets the consequentialist criterion of right action. When they are not, when bias, opacity, or dehumanisation dominate, AI recruitment fails the ethical test and demands corrective governance. This study reinforces the insight of Segall that justice and equality are indispensable components of consequentialist morality. Efficiency without fairness does not constitute moral progress. For AI recruitment to align with South Africa's constitutional commitment to dignity and equality, technological advancement must be embedded in ethical oversight, transparency, and social accountability.

In conclusion, the integration of consequentialist ethics into AI governance provides both a diagnostic and prescriptive tool. Diagnostically, it reveals where AI recruitment generates harmful consequences that contradict social welfare. Prescriptively, it guides policymakers and practitioners toward designing systems that maximise positive outcomes for all stakeholders. This ethical orientation transforms consequentialism from an abstract theory into a practical instrument for ensuring that digital transformation advances not only economic productivity but also moral and social justice.

6.3 Recommendations

Drawing from the findings and the consequentialist evaluation of AI recruitment, this section presents a set of actionable recommendations designed to ensure that the implementation of

Artificial Intelligence (AI) in South Africa's recruitment processes maximises collective welfare and minimises social harm. Consequentialism provides the normative basis for these recommendations: each proposed measure aims to promote outcomes that enhance fairness, equity, and human dignity while preventing or mitigating potential ethical risks such as bias, opacity, and exclusion. The recommendations are therefore structured across four interdependent domains: legal and policy, organisational and governance, educational and capacity development, and research and innovation, ensuring a holistic response to the ethical issues identified throughout the study.

6.3.1 Safeguard Human Dignity in Recruitment

Consequentialist ethics asserts that the ultimate moral measure of any system lies in its tangible impact on human welfare and dignity. Accordingly, AI recruitment processes must be designed to uphold the intrinsic worth of every candidate, ensuring that individuals are treated not merely as data points but as autonomous moral agents deserving of respect. The use of AI in hiring should, therefore, prioritise human well-being, fairness, and inclusivity, key outcomes that signify positive moral consequences under the consequentialist framework. To achieve this, recruitment systems should incorporate mechanisms that preserve human interaction and provide opportunities for candidates to engage meaningfully with the process. Applicants must receive clear explanations for hiring decisions and be afforded the right to appeal or contest automated rejections (Newlands, 2021, p. 92). Transparent communication throughout the recruitment journey mitigates the psychological harm often associated with dehumanisation and aligns with Section 10 of the Constitution of the Republic of South Africa (1996), which guarantees the right to human dignity. Maintaining empathy, dialogue, and responsiveness in hiring practices ensures that technological efficiency does not eclipse moral responsibility.

Practically, organisations should implement human-in-the-loop models, where human recruiters review and validate AI-generated recommendations at key decision points. This hybrid approach balances technological efficiency with moral discernment, producing outcomes that are both ethically sound and operationally effective. Human oversight also strengthens trust between employers and applicants, fostering a culture of fairness and accountability that enhances the overall welfare of the recruitment ecosystem. Beyond the organisational level, policymakers must institutionalise the protection of dignity within national AI governance structures. While existing legislation, such as the Protection of Personal Information Act (POPIA, 2013) and the Employment Equity Act (EEA, 1998), addresses

privacy and discrimination, they fall short of regulating the ethical dimensions of algorithmic decision-making. As Raghavan et al. (2020, p. 7) observe, the absence of explicit mechanisms for transparency, contestability, and redress undermines procedural justice and erodes public confidence in automated systems.

To bridge this gap, the establishment of a National AI Ethics and Governance Framework, coordinated by the Department of Communications and Digital Technologies, is essential. This framework could mandate regular algorithmic audits, bias testing, and ethical certification of recruitment technologies to ensure compliance with fairness and explainability standards. A National AI Ethics Council should be empowered to oversee these functions, aligning South Africa with international benchmarks such as UNESCO's Recommendation on the Ethics of Artificial Intelligence (2021) and the World Economic Forum's (2025) global principles for responsible AI governance. Legislative reform should further enshrine a "right to explanation" in AI-based recruitment, requiring employers to justify automated outcomes. It should also introduce mandatory impact assessments before AI deployment and impose penalties for discriminatory or opaque AI practices. These legal and institutional safeguards would ensure that the consequences of AI recruitment fairness, equity, and human welfare are not merely aspirational ideals but enforceable rights. In doing so, South Africa can exemplify how technological advancement and ethical integrity can coexist within a consequentialist moral framework.

6.3.2 Institutionalise Transparency and Human Oversight

Transparency and human oversight are indispensable pillars of ethical AI governance, particularly in recruitment contexts where algorithmic decisions directly affect people's livelihoods and dignity. From a consequentialist standpoint, transparency and accountability hold moral value because they produce beneficial social outcomes such as trust, fairness, and legitimacy while preventing harms such as bias, opacity, or exclusion. Consequently, decision-making authority in recruitment cannot be entirely delegated to autonomous systems; moral responsibility for outcomes must remain with human agents. Organisations should therefore embed transparency mechanisms and human accountability protocols into all AI-assisted recruitment processes. This entails ensuring that every stage of decision-making, from candidate screening to final selection, can be explained, reviewed, and justified. As Selbst and Barocas (2018, p. 59) argue, maintaining a human-in-the-loop is essential for preserving ethical agency: humans must retain the capacity to question, interpret, and, where appropriate,

overturn algorithmic outputs. Such oversight not only prevents unjust outcomes but also fulfils the consequentialist imperative to maximise welfare and minimise harm. Transparency must also extend outward to candidates and the public. Firms should disclose the use of AI tools in recruitment and provide clear, accessible explanations of how these systems assess applications. Candidates have a moral right, consistent with South Africa's constitutional principles of fairness and dignity, to understand how decisions affecting their employment opportunities are made. Open communication builds trust, strengthens institutional legitimacy, and contributes to positive moral outcomes such as empowerment and inclusion.

6.3.3 Moderate Algorithmic Bias through Inclusive Data and Regular Audits

Algorithmic bias represents one of the most profound ethical challenges in the deployment of Artificial Intelligence (AI) recruitment systems. When algorithms reproduce or amplify historical inequalities embedded in training data, they perpetuate patterns of exclusion that directly harm social welfare (Bolukbasi et al., 2019, p. 8). From a consequentialist standpoint, such outcomes are morally unacceptable because they generate harm and inequity rather than collective benefit. Any recruitment system that systematically disadvantages certain groups, whether by race, gender, language, or socio-economic background, fails the consequentialist test of maximising good and minimising harm. Therefore, employers, developers, and policymakers bear an ethical duty to actively identify, mitigate, and prevent algorithmic bias through inclusive data governance and rigorous auditing practices. To achieve these objectives, training datasets must reflect the demographic and cultural diversity of South Africa's population. This includes recognising and ethically representing variables such as race, gender, disability, language, and educational background. Data collection should be guided by fairness obligations enshrined in the Employment Equity Act (1998), which mandates the elimination of barriers to equitable representation. AI systems that rely on homogenous or incomplete datasets risk entrenching systemic inequalities that the South African Constitution seeks to dismantle. Inclusive data practices, therefore, serve not only a technical purpose but a moral one: they ensure that the consequences of AI recruitment align with the country's constitutional commitment to justice and equality.

Regular algorithmic audits must be institutionalised as part of ongoing ethical compliance programmes. These audits should systematically assess whether AI recruitment systems disproportionately disadvantage historically marginalised groups or produce unintended discriminatory outcomes. According to Hunkenschroer and Luetge (2022, p. 985), continuous

auditing enhances accountability by enabling the early detection of algorithmic bias and promoting transparency in decision-making processes. To ensure impartiality and credibility, such audits should include independent third-party reviews as well as internal fairness assessments. This dual-layered oversight strengthens organisational accountability and builds public trust in AI-driven hiring practices. Consequentialist ethics strongly support such proactive interventions because they directly improve welfare outcomes by enhancing fairness, inclusion, and access to opportunity. When algorithmic bias is identified and corrected, AI systems do not merely gain technical precision; they acquire moral legitimacy by promoting justice and reducing harm. As Bogen and Rieke (2020, p. 15) and Zou and Schiebinger (2021, p. 324) demonstrate, biased data inevitably lead to exclusionary results. By contrast, de-biased and representative datasets foster equitable recruitment outcomes that benefit both individuals and society, thereby satisfying the consequentialist imperative to maximise overall well-being.

In practical terms, AI systems must be trained on representative, context-sensitive datasets that capture South Africa's linguistic, cultural, and educational diversity. This includes recognising qualifications from historically Black universities, experiences gained in informal sectors, and multilingual resumes that reflect the realities of South African society. To institutionalise these practices, government regulators and technology providers should co-develop national data-equity standards, requiring dataset documentation, demographic parity testing, and recurring fairness audits. These standards should be embedded within a broader AI Governance Framework overseen by the Department of Communications and Digital Technologies and aligned with international benchmarks such as UNESCO's Recommendation on the Ethics of Artificial Intelligence (2021). Furthermore, data diversity initiatives should involve multi-stakeholder partnerships with universities, non-governmental organisations, and public-sector bodies to ensure the ethical inclusion of underrepresented communities in algorithmic design. As the consequentialist principal dictates, ethical foresight, anticipating and preventing harm, is superior to reactive correction after damage has occurred. By embedding fairness, inclusivity, and continuous auditing into AI recruitment systems, South Africa can ensure that technological progress serves the moral purpose of advancing justice, equality, and collective welfare.

6.3.4 Build Ethical Capacity and Awareness

Sustainable ethical governance depends fundamentally on the cultivation of ethical literacy, moral reasoning, and technological competence among all stakeholders involved in AI

recruitment. Consequentialist ethics emphasises that the moral value of any action, including technological innovation, must be evaluated in terms of its real-world outcomes. Therefore, building ethical capacity ensures that decision-makers across institutions are equipped not only to understand AI systems but to assess their societal consequences critically. Educational institutions have a significant role in embedding ethical competence at the foundation of professional formation. Universities and training bodies should integrate AI ethics, digital justice, and social impact courses into curricula for business management, computer science, and human resources programmes. These courses must go beyond abstract moral theory to include applied training on ethical decision-making frameworks, particularly consequentialist reasoning, which evaluates technology according to its effects on welfare, fairness, and inclusion. As Marwala (2024, p. 71) asserts, capacity building in digital ethics is a prerequisite for Africa's equitable participation in the fourth industrial revolution, ensuring that technological progress serves as a catalyst for justice rather than a source of division.

At the organisational level, regular ethics and accountability training should be institutionalised for HR practitioners, data scientists, and AI developers. Such training would cultivate moral awareness of how data inputs, design choices, and algorithmic outputs affect candidates' lives and societal welfare. Organisations should also develop cross-disciplinary ethical learning programmes that encourage collaboration between ethicists, technologists, and policymakers, fostering shared responsibility for the moral quality of AI systems. Public and civil society initiatives can further strengthen ethical awareness by creating spaces for dialogue on the social consequences of automation and digital governance. This participatory approach ensures that AI ethics does not remain confined to academic or corporate elites but becomes a shared societal project grounded in inclusivity and transparency.

6.3.5 Reduce Over-Reliance through Governance and Accountability Frameworks

Consequentialist ethics demands that technology serve human welfare rather than replace human judgment. Over-reliance on Artificial Intelligence (AI) in recruitment undermines this moral balance by delegating critical decisions about people's livelihoods to systems that lack empathy, contextual understanding, and moral reasoning. While AI offers efficiency and analytical precision, it must operate within clear ethical and institutional boundaries that preserve human oversight and accountability. According to Driver (2015, pp. 97–98), moral rightness depends on the overall consequences of an action; thus, decision-making systems must be designed to maximise welfare and minimise harm. If AI-driven recruitment leads to

alienation, opacity, or moral disengagement, its use becomes ethically unjustifiable under consequentialist reasoning. To counteract the risks of over-reliance, organisations must embed human governance structures within every stage of the recruitment process. Human resources (HR) professionals should remain active participants reviewing AI outputs, contextualising candidate assessments, and ensuring that decisions align with ethical and legal standards. This hybrid model, often referred to as “human-in-the-loop” governance, ensures that technological tools support rather than supplant moral judgment (Selbst and Barocas, 2018, p. 59). By maintaining human oversight, organisations preserve moral agency and safeguard against harmful or unintended consequences such as wrongful exclusion, biased filtering, or procedural unfairness.

A robust governance and accountability framework should also include transparent documentation of AI decision-making processes, risk assessments, and ethical impact reports. Each system’s design, purpose, and limitations must be clearly defined and subject to regular review. This aligns with the consequentialist imperative of evaluating outcomes continually to ensure that benefits continue to outweigh harms. Governance structures should mandate periodic reviews of AI performance metrics, particularly false positives, and false negatives in candidate selection, to ensure fairness and accuracy over time. As Marwala (2024, p. 71) argues, ethical governance and capacity-building are basics for Africa’s equitable participation in the Fourth Industrial Revolution. At the institutional level, organisations should establish Ethical AI Oversight Committees composed of HR experts, ethicists, data scientists, and legal advisors. These committees would be responsible for monitoring compliance with both internal codes of conduct and external legal frameworks such as the Employment Equity Act (1998), Protection of Personal Information Act (POPIA, 2013), and Promotion of Administrative Justice Act (PAJA, 2000). Integrating these governance mechanisms ensures that recruitment technologies operate transparently and within the boundaries of South Africa’s constitutional commitment to fairness, dignity, and equality.

On a national scale, the Department of Communications and Digital Technologies, in collaboration with the Information Regulator and the Commission for Employment Equity, should spearhead the creation of a National AI Governance and Accountability Framework. This framework would formalise mandatory reporting standards, ethical certification for recruitment tools, and mechanisms for redress in cases of algorithmic discrimination or procedural injustice. Aligning with international benchmarks such as UNESCO’s Recommendation on the Ethics of Artificial Intelligence (2021) and the World Economic

Forum's AI Governance Principles (2025) would further strengthen South Africa's leadership in responsible innovation. From a consequentialist perspective, reducing over-reliance on AI through robust governance produces outcomes that maximise social good. It restores human judgment where moral sensitivity is required, prevents systemic harms from automated decision-making, and ensures that technological efficiency does not eclipse ethical responsibility. By embedding accountability, transparency, and human oversight into recruitment governance, organisations and policymakers uphold the moral principle that technology must serve people, not the reverse. This balance ensures that AI-driven recruitment contributes to societal welfare, justice, and the collective flourishing of all stakeholders in the employment ecosystem.

In closing, this study affirms that technology, while transformative, is never morally neutral. The ethical value of Artificial Intelligence in recruitment and selection lies not in its computational precision but in the human consequences it produces. Through the consequentialist lens, morality is measured not by intention but by impact, by whether the outcomes of technological innovation enhance human welfare, justice, and dignity. South Africa's pursuit of digital transformation must therefore be guided by the moral conviction that progress is ethical only when it uplifts the collective good. AI recruitment systems, when designed and governed responsibly, can become instruments of empowerment, bridging opportunity gaps, fostering inclusion, and strengthening institutional trust. Yet, without ethical vigilance, they risk deepening inequality and eroding the very values that sustain a just society. The findings of this dissertation thus call for an enduring moral partnership between ethics and innovation, where every algorithm is held accountable to the lives it affects. The study concludes that the true measure of technological advancement is not efficiency alone, but the preservation of human dignity within systems of progress. Consequentialism offers a guiding compass for this journey: to act, design, and govern in ways that maximise welfare, minimise harm, and ensure that every digital step forward remains a moral stride toward justice.

References

- Abrahams, L. et al. (2023) 'Decoding Bias in AI Video Interviews', *South African Journal of Psychology*, 53(1), pp. 1–12.
- Acemoglu, D. and Restrepo, P. (2019), The wrong kind of AI? Artificial Intelligence and the future of labour demand, *Journal of Economic Perspectives*, 33(2), pp. 51–72.
- Adams, R. (2024). *Artificial Intelligence and the Public Interest in Africa: Towards Inclusive Governance Frameworks*. Johannesburg: HSRC Press.
- Ajunwa, I., 2020. Algorithmic bias and discrimination in AI recruitment. *Harvard Law Review*, 12(3), pp.12–13.
- Ajunwa, I. (2020) 'The paradox of automation as anti-bias intervention', *Cardozo Law Review*, 41(5), pp. 1671–1745.
- Arrhenius, G., Bykvist, K. and Campbell, T., 2020. Population ethics and intergenerational justice. *Ethics and Philosophy*, 45(2), pp.115–120.
- Asoba, S.N. and Mefi, N.P., 2021. Emerging talent management themes in HRM 4.0: A study in the banking sector in South Africa. *Academy of Strategic Management Journal*, 20(6), pp.1-6.
- Asoba, Samson, and Mefi, Nteboheng. (2025). *The Impact of Artificial Intelligence Adoption on Unethical Recruitment and Selection Practices in the South African Public Sector*. Available at: [Researchgate](#) (Accessed: 22 October 2025).
- Barocas, S., and Selbst, A. D. (2016). Big Data's Disparate Impact. *California Law Review*, 104(3), 671–732.
- Bentham, J., 1789. *An Introduction to the Principles of Morals and Legislation*. London: T. Payne and Son.
- Binns, R. (2020) 'Algorithmic accountability and public reason', *Philosophy and Technology*, 33(4), pp. 519–565.
- Binns, R. (2018). *Fairness in Machine Learning: Lessons from Political Philosophy*. *Proceedings of the 2018 Conference on Fairness, Accountability, and Transparency*, 81, pp.45–60.
- Birhane, A. (2020), Algorithmic colonization of Africa, *SCRIPTed*, 17(2), pp. 389–409.

- Blair, C. (2025) 'Incorporating AI into workforce planning: Optimising skills for the South African workforce – Part 1'. EBNet. Available at: [EBNet](#) (Accessed: 26 September 2025).
- Bogen, M. and Rieke, A. (2018) Help Wanted: An Examination of Hiring Algorithms, Equity, and Bias. Washington, DC: Upturn. Available at: [Upturn](#) (Accessed: 26 September 2025).
- Bolukbasi, T. et al. (2019), Man is to Computer Programmer as Woman is to Homemaker? Proceedings of NeurIPS, pp. 1–15.
- Bowen, G. (2009) 'Document Analysis as a Qualitative Research Method', Qualitative Research Journal, 9(2), pp. 27–40.
- Braun, V., and Clarke, V. (2006) 'Using Thematic Analysis in Psychology', Qualitative Research in Psychology, 3(2), pp. 77–101.
- Buolamwini, J. and Gebru, T. (2018) 'Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification', Proceedings of Machine Learning Research, 81, pp. 1–15.
- Burrell, J. (2016) 'How the Machine "Thinks": Understanding Opacity in Machine Learning Algorithms', Big Data and Society, 3(1), pp. 1–12.
- Cappelli, P. et al. (2021) 'AI in Hiring: Assessing the Risks and Rewards', Harvard Business Review, 99(3), pp. 25–34.
- Card, D. and Smith, N.A. (2020). On consequentialism and fairness. Frontiers in Artificial Intelligence. Available at: [FRONTIERS](#) (Accessed: 30 September 2025).
- Chamorro-Premuzic, T. et al. (2022) 'The AI-Powered Workplace', Harvard Business Review, 100(2), pp. 12–20.
- Chamorro-Premuzic, T., Akhtar, R., Winsborough, D. and Sherman, R.A., 2020. The potential for AI in recruitment: Efficiency and fairness. Journal of Applied Psychology, 105(1), pp.70–82.
- Chander, A. (2021) 'The Racist Algorithm?', Michigan Law Review, 119(5), pp. 40–58.
- Chen, Z. (2023) 'Ethics and discrimination in AI-enabled recruitment practices', Humanities and Social Sciences Communications, 10(1), pp. 1–12.

- Citron, D., and Pasquale, F. (2014) 'The Scored Society: Due Process for Automated Predictions', *Washington Law Review*, 89(1), pp. 1–33.
- Cliffe Dekker Hofmeyr. (2025). How algorithmic bias in AI hurts your business and what you can do. Available at: [Cliffe Dekker Hofmeyr Publications](#). (Accessed: 30 September 2025).
- CMS Law. (2024). Bias in the code: How AI recruitment processes can land employers in hot water. Available at: [CMR LAW](#) (Accessed: 15 October 2025).
- Cowgill, B. (2019) 'Bias and Productivity in Humans and Algorithms', Columbia Business School Research Paper, pp. 1–30.
- Dastin, J. (2018) 'Amazon scraps secret AI recruiting tool that showed bias against women', Reuters. Available at: [Reuters](#) (Accessed: 26 September 2025).
- Datt, G. et al. (2021) 'AI Recruitment Technologies in Global South Contexts', *Journal of International HR Management*, 12(3), pp. 1–15.
- Datt, G. et al. (2022) 'AI Compliance in Emerging Markets: The South African Challenge', *International Journal of HR Technology*, 3(2), pp. 1–12.
- Datt, G. et al. (2022) 'Contextual AI for Global South Recruitment', *Journal of International HR Management*, 13(4), pp. 1–15.
- de Lazari-Radek, K. and Singer, P., 2017. *The Point of View of the Universe: Sidgwick and Contemporary Ethics*. Oxford: Oxford University Press.
- Department of Employment and Labour. (2023). *Employment equity and workplace inclusivity in South Africa*. Government of South Africa. Available at: [Department of Employment and Labour](#) (Accessed: 20 September 2025).
- Diakopoulos, N. (2016). Accountability in Algorithmic Decision Making. *Communications of the ACM*, 59(2), 56–62.
- Dlamini, N. (2024), Digital Divides in AI Recruitment, *Transformation*, 113, pp. 20–35.
- Dlamini, N. and Patel, L. (2023) 'AI Hiring in South Africa: Efficiency vs Transformation', *SA Journal of Industrial Psychology*, 49, pp. 1–9.
- Driver, J., 2015. Consequentialism and practical reasoning. *Philosophy Compass*, 10(2), pp.95–105.

- Edwards, L., and Veale, M. (2018) 'Enslaving the Algorithm: From a "Right to an Explanation" to a "Right to Better Decisions"?', IEEE Security and Privacy, 16(3), pp. 46–54.
- IEEE. (2021). Ethically Aligned Design: A Vision for Prioritizing Human Well-being with Autonomous and Intelligent Systems. IEEE Global Initiative on Ethics of Autonomous and Intelligent Systems, pp. 1-139.
- Fabris, A., Ndlovu, P. and Visser, T. (2025) 'Fairness and bias in algorithmic hiring: A multidisciplinary perspective', ACM Transactions on Internet Technology, 25(2), pp. 1–24.
- Fritts, B. and Cabrera, J. (2021) 'Automation and the human experience in hiring', Ethics and Information Technology, 23(4), pp. 791–801.
- Floridi, L., and Cowls, J. (2019). A Unified Framework of Five Principles for AI in Society. Harvard Data Science Review, pp. 1-16.
- Giwa, A., Mthembu, L. and Daniels, K. (2024) 'Algorithmic bias and recruitment inequities in African labour markets', African Journal of Information Systems, 16(1), pp. 1–15.
- Graylink. (2024) 'Local vs Global: The Case for South African ATS Software'. Available at: [Graylink](#) (Accessed: 25 September 2025).
- Graylink. (2025) 'WhatsApp recruiting: The new frontier in candidate engagement', Graylink Blog, 22 July. Available at: [Graylink](#) (Accessed: 24 September 2025).
- Hare, R.M., 1981. Moral Thinking: Its Levels, Method, and Point. Oxford: Clarendon Press, pp. 229-236.
- Harwell, D. (2019) 'A face-scanning algorithm increasingly decides whether you deserve the job', The Washington Post, 22 October. Available at: [Washington Post](#) (Accessed: 10 September 2025).
- HireSales.ai. (2025). *AI-powered recruitment for sales talent*. Available at: [HireSales](#) (Accessed: 10 September 2025).
- HR Personnel Services. (2025). The Ethics of AI in Recruitment: Balancing Efficiency and Fairness. Available at: [hrpersonnelservices](#) (Accessed: 20 October 2025).

- Hunkenschroer, A., and Luetge, C. (2022). Ethical AI in human resources: Consequences and accountability. *AI and Society*, 37(3), pp. 971–985.
- Innova Solutions. (2025), From Automation to Accountability: Embracing Ethical AI in Workforce Recruitment. Available at: [Innova](#) (Accessed: 22 October 2025).
- Institute for Security Studies (ISS) (2023) *Procedural justice and policing: Building trust in South Africa's police (Southern Africa Report 54)*. Pretoria: ISS.
- International Monetary Fund (IMF) (2020) 'Six Charts Explain South Africa's Inequality'. Available at: [IMF](#) (Accessed: 15 September 2025).
- Keller, P. and Aplin, T. (2024) 'Trade secrets and AI transparency: Challenges in disclosure regimes', *European Intellectual Property Review*, 46(1), pp. 9–11.
- Köchling, A. and Wehner, M. (2020) 'Discriminated by an Algorithm', *Computers in Human Behavior*, 109, 106339.
- Laher, S. (2022) 'Cross-Cultural Psychometrics in AI Assessments', *SA Journal of Industrial Psychology*, 48, pp. 1–8.
- Labour Court of South Africa (2021) *Moyo v. Commission for Conciliation, Mediation and Arbitration*, Case No. JR 1022/19.
- Limarga, A., Sutoyo, J., and Santoso, H. (2023). Formalising consequentialist principles for machine ethics: Evaluating AI systems in social contexts. *Journal of Artificial Intelligence and Ethics*, 7(2), pp. 440–448. [ijcai.org](#) (Accessed: 13 September 2025).
- Lund, S. et al. (2022). *The State of AI in Talent Acquisition*. McKinsey Global Institute, pp. 1-21.
- Maharaj, C. and Obalade, A. (2025) 'Artificial Intelligence innovations in talent recruitment, retention and diversity mapping within South Africa: A meta-narrative analysis', *International Journal of Innovation and Technology Management*, 6(1), pp. 1–15.
- Maimela, C. (2025). Artificial Intelligence in South African universities: curriculum, PMC. Available at: [Frontiers](#) (Accessed: 22 October 2025).
- Maleka, M., Townes, J., and Marange, T. (2025). Recruitment in Organisations: A Qualitative Perspective from A Developing Country, pp. 1-22.

- Mapungwana, P. and Chadyiwa, M. (2025). Synthesising the potential of Artificial Intelligence in the fulfilment of sustainable development goals in South Africa: An ethical perspective. *ScienceDirect*, pp. 1-10.
- Marwala, T. (2024) 'Inclusivity Must Be at the Heart of AI Development: Address at the ICEGOV 2024 Conference'. United Nations University Operating Unit on Policy-Driven Electronic Governance (UNU-EGOV) / University of the Western Cape, Cape Town, Available at: [WSG](#) (Accessed: 26 September 2025).
- McKinsey and Company (2024), Building AI Trust: The Key Role of Explainability. *QuantumBlack, AI by McKinsey*. Available at: [McKinsey](#) (Accessed: 6 September 2025).
- Mejía, R. (2019) 'Ethics of consequentialism in technological systems', *Journal of Applied Philosophy*, 36(3), pp. 409–421.
- Mendis, R., 2024. A global outlook on 13 AI laws affecting hiring and recruitment. HR Executive. Available at: [Hrexecutive.com](#) [Accessed 20 October 2025].
- Mhlongo, T., and Mabuza, L. (2022). Reducing bias in recruitment: Opportunities and challenges of AI in South Africa. *Journal of Human Resource Management*, 20(2), pp. 45–58.
- Mill, J.S., 1879. Utilitarianism. London: Parker, Son, and Bourn, pp. 1-121.
- Mkhize, S. (2023) 'Localizing AI for African Labor Markets', *African Journal of Management*, 11(2), pp. 110–125.
- Mobley v. Workday, Inc., No. 3:23-cv-00770 (N.D. Cal. July 12, 2024). Available at: [Seyfarth](#) (Accessed: 30 July 2025).
- Mokgoro, Y., 2023. AI ethics and regulatory compliance in South Africa. *South African Journal of Law and Technology*, 29(1), pp.1–15.
- Mondo. (2025). The Ethics of AI in Recruitment: Balancing Efficiency and Equality. Available at: [MONDO](#) (Accessed: 19 October 2025).
- MTN South Africa (2023) 'MTN South Africa takes new AI step as it launches chatbot "SiYa"', *Engineering News*, 4 December. Available at: [Engineering News](#) (Accessed: 18 July 2025).
- MyBroadband (2021) 'Vodacom, MTN and Telkom WhatsApp chatbots compared', 7 December. Available at: [MyBroadband](#) (Accessed: 24 September 2025).

- Mylly, T. (2023) 'Intellectual property and access to information in the age of AI', *European Law Journal*, 29(6), pp. 1005–1024.
- Ncube, T.R. (2025). The impact of Artificial Intelligence on human resource management practices: An investigation, SAJHRM. Available at: scielo.org.za (Accessed: 24 October 2025).
- Newlands, G., 2021. Dehumanization and algorithmic hiring. *Ethics and Information Technology*, 23(2), pp.90–95.
- Nkosi, B. (2023) 'Algorithmic Affirmative Action? AI and Employment Equity in South Africa', *Transformation*, 111, pp. 110–128.
- Nkosi, B. and Ruggunan, S. (2022) 'Digital Hiring and Employment Equity in Post-Apartheid South Africa', *Transformation*, 108, pp. 1–18.
- Nkosi, B. and Ruggunan, S. (2023) 'Decolonizing AI Recruitment: A South African Perspective', *Transformation*, 110, pp. 1–20.
- Nkosi, M. and Ruggunan, S. (2023) 'Ethics, equity, and Artificial Intelligence in South African workplaces', *Transformation in Higher Education*, 8(1), pp. 1–20.
- Nzimakwe, T. and Reddy, P. (2022) 'AI and Employment Equity in South Africa', *SA Journal of Human Resource Management*, 20, pp. 1–12.
- Obasa, A.E. (2023). The ethics of Artificial Intelligence in healthcare settings. Stellenbosch University. Available at: scholar.sun.ac.za (Accessed: 22 October 2025).
- OCEG (2023) *Transparency and Risk in Corporate Governance: AI Case Studies*. Phoenix, AZ: OCEG. Available at: [OCEG](https://www.oci.gov) (Accessed: 19 September 2025).
- Parfit, D., 1984. *Reasons and Persons*. Oxford: Clarendon Press, pp. 1-156.
- Pasipamire, N. (2024). Navigating algorithm bias in AI: ensuring fairness and trust in Africa PMC. Available at: [pmc.ncbi.nlm.nih.gov](https://pubmed.ncbi.nlm.nih.gov) (Accessed: 22 October 2025).
- Raghavan, M., Barocas, S., Kleinberg, J. and Levy, K. (2020) 'Mitigating bias in algorithmic hiring: Evaluating claims and practices', in *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency (FAccT '20)*, pp. 469–481.
- Republic of South Africa (1998) *Employment Equity Act 55 of 1998*. Pretoria: Government Printer.

- Republic of South Africa (2013) *Protection of Personal Information Act 4 of 2013*. Pretoria: Government Printer.
- Republic of South Africa (2024) *National Artificial Intelligence Policy Framework*. Pretoria: Department of Communications and Digital Technologies.
- Rigotti, C. (2024), Fairness, AI and recruitment: A scoping review, *Computer Law and Security Review*, 52, 105923.
- Ringover (2024), What Is Mobile Recruiting? 6 Effective Strategies. Available at: [Ringover](#) (Accessed: 12 July 2025).
- Rodriguez, V. (2023) 'Ethical Implications of AI-Based Algorithms in Recruiting'. *Master's Graduate Project*, Boise State University. Available at: [Boise State](#) (Accessed: 19 September 2025).
- Roos, A. (2021) 'POPIA and AI: South Africa's Balancing Act', *SA Journal of Information Law*, 19(2), pp. 110–125.
- Ruttkamp-Bloem, E. (2021) 'AI technology has, from its inception, been meant to augment what it means to be human...', *South African Journal of Philosophy*, 40(3), pp. 1–3.
- Saunders, M. et al. (2018), *Research Methods for Business Students*. 7th end. Harlow: Pearson, pp. 1-70.
- Segall, S., 2021. Consequentialism and structural injustice. *Philosophy and Public Affairs*, 49(2), pp.120–140
- Segall, S. (2021) 'Inequality, innovation, and the ethics of AI', *Ethics and International Affairs*, 35(3), pp. 359–366.
- Selbst, A.D., and Barocas, S. (2018) 'The intuitive appeal of explainable machines', *Fordham Law Review*, 87(3), pp. 1085–1139.
- SHRM LABS (2025) 'Fresh SHRM Research Explores Use of Automation and AI in HR'. Available at: [SHRM](#) (Accessed: 20 September 2025).
- Silverman, D. (2020), *Interpreting Qualitative Data*. 6th end. London: Sage, pp. 1-13.
- Singer, P. (2015), *The Most Good You Can Do: How Effective Altruism is Changing Ideas About Living Ethically*. New Haven: Yale University Press, pp. 40-232.

- South African Business Matters (2025), Incorporating AI into Workforce Planning: Optimising Skills for the South African Workforce of the Future, Part 1. Available at: [Business Matters](#) (Accessed: 10 July 2025).
- Statistics South Africa (2023) Quarterly Labour Force Survey Q4 2023. Pretoria: Statistics South Africa.
- Statistics South Africa (2025) Quarterly Labour Force Survey Q2 2025. Pretoria: Statistics South Africa.
- Tacit Base. (2025). Making AI Accountable: The Role of Governance and Transparency in Recruitment. Available at: tacitbase.com (Accessed: 22 October 2025).
- Tambe, P., Cappelli, P. and Yakubovich, V., 2019. Artificial Intelligence in human resources management: Challenges and a path forward. *California Management Review*, 61(4), pp.15–42.
- Tambe, P. et al. (2022) 'AI and Human Capital Management', *Journal of Labor Economics*, 40(S1), pp. 293–324.
- TechFinancials (2025) 'Khumbula Tech Trying to Take the Lead in AI ethics in SA'. Available at: [TechFinancials](#) (Accessed: 25 September 2025).
- UNICEF Innovation Fund (2021), Giraffe: Job matching for youth employment in South Africa. Available at: [UNICEF](#) (Accessed: 26 September 2025).
- Upadhyay, A. and Khandelwal, K. (2020) 'AI in Recruitment: A Conceptual Framework', *Journal of Business Research*, 115, pp. 10–25.
- Van Esch, P., Black, J.S. and Ferolie, J. (2021) 'Marketing AI recruitment: The next phase in job application and selection', *Computers in Human Behavior*, 114, 106586.
- Van Esch, P. et al. (2021) 'AI Recruitment and Candidate Perceptions', *International Journal of Selection and Assessment*, 29(1), pp. 1–15.
- Villasenor, J. (2024) 'Trade secrecy and transparency in AI regulation', *Brookings Tech Policy Series*, pp. 45–47.
- Vivek, R. (2022), Is Blind Recruitment an Effective Recruitment Method?, *International Journal of Applied Research in Business and Management*, 3(3). Available at: [IJARBM](#) (Accessed: 30 July 2025).

- Wachter, S. et al. (2018) 'Why a Right to Explanation of Automated Decision-Making Does Not Exist', *International Data Privacy Law*, 7(2), pp. 20–35.
- World Bank (2022), Gini Index (World Bank estimate), South Africa. Available at: [World Bank](#) (Accessed: 25 September 2025).
- World Economic Forum (2025), Gender Parity in the Intelligent Age. Available at: [WEF](#) (Accessed: 29 July 2025).
- Yılmaz, O., Singh, P., and Choudhury, A. (2018) 'Multilingual natural language processing challenges in recruitment', *Computational Linguistics Review*, 29(2), pp. 55–60.
- Zarsky, T. (2021) 'Privacy and Manipulation in the Digital Age', *Theoretical Inquiries in Law*, 22(1), pp. 1–36.
- Zielinski, D. (2020) 'Artificial Intelligence in the Workplace', *Society for Human Resource Management (SHRM)*. Available at: [SHRM](#) (Accessed: 16 September 2025).
- Zou, J. and Schiebinger, L., 2021. AI can be sexist and racist. It is time to make it fair. *Nat's*, 559(7714), pp.324–326.