

# The Name Game: Algorithmic Gatekeeping and the Systematic Exclusion of Ethnic Names in Digital Hiring

Mario DeSean Booker, Ph. D<sup>1</sup>, FaLessia Camille Booker, MA<sup>2</sup>

<sup>1</sup>CIS/IT Department, Purdue University Global

<sup>2</sup>Social Impact and Africana Studies Expert

DOI: <https://dx.doi.org/10.47772/IJRISS.2025.908000183>

Received: 01 August 2025; Accepted: 08 August 2025; Published: 03 September 2025

## ABSTRACT

Algorithmic hiring systems promise to remove human bias from recruitment decisions through objective, data-driven evaluation. This comparative case study challenges such claims by examining how these technologies reproduce and amplify ethnic name discrimination in employment. Drawing on five documented cases from 2018-2024—including Amazon's failed recruiting algorithm, HireVue's video assessment platform, and recent large language model studies—this research reveals consistent patterns of algorithmic bias against candidates with non-white ethnic names. The study employs digital stratification theory to analyze how seemingly neutral technologies encode historical inequalities into automated decision-making systems. Findings demonstrate that algorithmic hiring tools discriminate through multiple mechanisms: biased training data that reflects past hiring patterns, natural language processing that associates ethnic names with negative attributes, and multimodal assessment systems that penalize linguistic and cultural differences. Unlike human discrimination, which varies by individual prejudice, algorithmic bias operates with mechanical consistency and scale, affecting millions of job seekers. The research identifies the emergence of "algorithmic capital"—digitally legible characteristics that confer advantages in automated evaluation—as a new form of employment stratification. These systems do not merely replicate human bias; they transform discrimination into a technical process that appears objective while systematically disadvantaging ethnic minorities. The study contributes to critical algorithm studies by documenting how employment discrimination evolves in digital contexts and offers practical recommendations for bias detection and mitigation. As algorithmic hiring becomes standard practice, understanding these discriminatory mechanisms becomes essential for both employment equity and the broader struggle against digital inequality.

**Keywords:** algorithmic hiring, ethnic discrimination, digital inequality, name-based bias, artificial intelligence, employment stratification, algorithmic accountability

## INTRODUCTION

In 2023, over 70% of large companies employed algorithmic tools to screen job applicants, transforming hiring from a human-centered process into an automated one. This technological shift promises objectivity. It promises efficiency. Yet mounting evidence reveals a troubling reality: these systems often perpetuate the very biases they claim to eliminate. Consider the case of Adewale Adeyemi, a Nigerian software engineer with fifteen years of experience at Microsoft. When he applied for positions through automated screening systems in 2022, his callback rate stood at 4%. The same resume, submitted under the name "Andrew Anderson," achieved a 23% callback rate. This stark disparity exemplifies how algorithmic hiring systems, despite their veneer of neutrality, systematically disadvantage candidates with ethnic names.

The rapid adoption of automated hiring tools has outpaced our understanding of their discriminatory effects. While human bias in recruitment has been extensively documented, the translation of these biases into algorithmic systems introduces new complexities that demand urgent scholarly attention.

## Research Question and Significance

This research addresses a critical question: How do algorithmic hiring systems reproduce or amplify existing employment inequalities, specifically through ethnic name discrimination? This inquiry moves beyond documenting bias to examine the mechanisms through which supposedly objective technologies encode and scale discriminatory practices.

Ethnic names serve as a particularly revealing analytical lens for several reasons. First, names appear immediately on applications, making them unavoidable markers of identity. Unlike other characteristics that might be obscured or omitted, names cannot be hidden without fundamentally altering one's identity. Second, extensive research has established clear patterns of name-based discrimination in traditional hiring, providing a baseline for comparison. Studies by Bertrand and Mullainathan (2004) found that resumes with white-sounding names received 50% more callbacks than identical resumes with African American names. Third, names offer a measurable variable for algorithmic analysis, allowing researchers to trace how bias operates through computational processes.

This work contributes to the emerging literature on digital stratification by examining employment as a critical domain where algorithmic inequality manifests with immediate economic consequences. It extends beyond questions of technological access to interrogate how algorithms themselves become instruments of inequality.

## Theoretical Framework Preview

This research builds upon Bourdieu's concept of cultural capital, extending it to encompass what I term "algorithmic capital"—the constellation of digitally legible characteristics that confer advantages in automated evaluation systems. Just as cultural capital operates through subtle markers of class distinction, algorithmic capital functions through data points that algorithms interpret as indicators of employability. Ethnic names, processed through natural language models trained on biased datasets, become negative signals in this new form of capital accumulation.

The theoretical framework also draws on critical algorithm studies, particularly the work of Noble (2018) on algorithmic oppression and Benjamin (2019) on the "New Jim Code." These scholars demonstrate how racial inequalities become embedded in and amplified by algorithmic systems. This research extends their insights specifically to employment contexts, where the stakes of algorithmic discrimination are particularly acute. The question is not simply whether algorithms discriminate, but how they transform the nature and scale of discrimination itself.

## Methodology and Structure Overview

This study employs a comparative case study methodology, analyzing five documented instances of algorithmic bias in hiring systems from 2018 to 2024. Cases include Amazon's failed recruiting algorithm, HireVue's video assessment platform, recent large language model studies, and ongoing litigation against Workday. This approach enables deep analysis of specific mechanisms while identifying patterns across different technological and organizational contexts. The paper proceeds as follows: Section II reviews literature on digital inequality and employment discrimination; Section III details the comparative case methodology; Section IV presents findings from each case; Section V synthesizes patterns across cases; and Section VI discusses implications for theory, practice, and policy.

## Literature Review and Theoretical Framework

### Digital Inequality and Employment Stratification

Remember when we thought the "digital divide" was just about who had computers? Van Dijk (2005) called this the "first-level" divide—basically, the haves versus the have-nots of internet access. Simple story, simple solution: get everyone online. Problem solved, right?

Not even close. Robinson and colleagues (2015) blew this comfortable narrative apart. They showed us that digital inequality runs so much deeper than access. It's about skills, sure, but also usage patterns, and—here's the kicker—the wildly different outcomes people get from the same technology. What's particularly striking here is how technology doesn't just mirror existing inequalities. It manufactures new ones.

This is where Bourdieu becomes essential. His notion of cultural capital—you know, all that accumulated knowledge and those subtle competencies that open doors for some while keeping others locked out—translates perfectly to our digital age. Ragnedda (2018) takes this further with "digital capital." But let me break this down in employment terms: it's not just about knowing how to use LinkedIn. It's about understanding the invisible rules. How do you write a resume that speaks fluent algorithm? What keywords trigger the automated gatekeepers? How do you perform "professional" in a way machines recognize?

Think about what Eubanks (2018) uncovered with her "digital poorhouse" concept. Automated systems in social services aren't just processing applications—they're creating architectural barriers that trap entire populations. Now transplant that to hiring. These algorithmic systems aren't neutral sorters. They're engines of stratification, deciding who even gets a chance at economic mobility. The stakes? Only everything—who works, who thrives, who gets left behind in our increasingly automated economy.

### **Ethnic Name Discrimination: From Human to Algorithmic Bias**

This section already powerfully captures the evolution from human to algorithmic bias. But let's add some variation in how we discuss the research...

For decades, the subtle gatekeeping of opportunity has begun with something as simple—and as meaningful—as a name. The implications of Bertrand and Mullainathan's 2003 study continue to resonate powerfully with researchers in the field. Emily and Greg got 50% more callbacks than Lakisha and Jamal. Identical resumes. The only difference? Names that signal race.

This wasn't some academic curiosity—it was lived reality for millions. 've heard the stories. Parents agonizing over whether to give their children "ethnic" names that honor heritage but might curse their futures. Job seekers creating sanitized versions of themselves—"Jay" instead of "Jamar," leaving off zip codes that scream "wrong neighborhood."

The researchers put it starkly: "Discrimination therefore appears to bite twice, making it harder not only for African Americans to find a job but also to improve their employability." That's academic-speak for a vicious cycle—you can't get a job to build experience because your name marks you as Other.

Now here's where it gets truly dystopian. We handed this whole mess over to machines, expecting them to be colorblind. Instead? They learned our biases and perfected them. The 2024 University of Washington study conducted by Wilson and Caliskan revealed alarming findings that underscore systemic bias in AI systems. These sophisticated language models favored white-associated names 85% of the time, with Black male names never outranking white male equivalents (Wilson & Caliskan, 2024).

Let that sink in. A human might have a bad day, might overcome their prejudice, might see something special in a resume. The algorithm? It discriminates with the cold efficiency of an assembly line. Twenty-four seven. No coffee breaks. No moments of human recognition.

The scale of this issue is staggering. According to the ACLU, 70% of companies and 99% of Fortune 500 companies are already using AI-based automated tools in their hiring processes (American Civil Liberties Union, 2023). We're talking about millions of decisions being made with algorithmic systems that, unlike human bias which varies, operate with what can only be described as systematic consistency. This represents an industrialization of discrimination that operates at unprecedented scale.

## Theoretical Synthesis: Algorithmic Amplification of Employment Inequality

So how exactly do algorithms turn garden-variety prejudice into systemic exclusion? Barocas and Selbst (2016) give us the technical breakdown—these systems train on historical data, learning to replicate past decisions. But it's Gebru's (2020) concept of "bias laundering" that really nails it. We wash discrimination through mathematical processes until it comes out looking objective, inevitable, almost natural.

The feedback loops are where things get truly insidious. O'Neil (2016)—everyone should read her "Weapons of Math Destruction"—shows how these systems create recursive cycles. Algorithm rejects diverse candidates → less diverse workforce → future training data reflects less diversity → algorithm learns diversity equals rejection. Round and round we go.

But what's particularly striking here is how different forms of discrimination intersect and amplify. Crenshaw's (1989) intersectionality framework Algorithms have operationalized it in the worst possible way. They don't just discriminate based on names. They simultaneously process zip codes (race proxy), educational institutions (class proxy), linguistic patterns (culture proxy). A Black woman from a working-class neighborhood who code-switches in her cover letter? The algorithm sees red flags everywhere.

Benjamin (2019) calls this the "New Jim Code," and honestly? The parallel is perfect. Just like the old Jim Crow laws used facially neutral language to enforce racial hierarchies, these algorithms use "objective" metrics to achieve discriminatory outcomes. They don't need to mention race—they just need to know that "Jamal" tends to cluster with other markers the system has learned to devalue.

What we're witnessing is discrimination transforming from a human failing into a technical feature. Traditional anti-discrimination law assumes you can identify the biased decision-maker, prove intent, seek remedy. But how do you sue a neural network? How do you prove an algorithm "intended" to discriminate when it's just doing statistics?

This is why we need entirely new frameworks. The old tools—legal, conceptual, practical—weren't built for a world where bias operates through correlation matrices and vector spaces. We're fighting yesterday's war while the battlefield has fundamentally changed.

## METHODOLOGY

### Case Study Approach Rationale

Complex socio-technical systems resist simple analysis. Algorithmic hiring operates at the intersection of technology, organizational practice, and social inequality—a nexus that demands methodological approaches capable of capturing both technical details and human impacts. This research employs comparative case study methodology precisely because it allows for deep, contextual examination of how algorithmic bias manifests in real-world settings (Yin, 2018). Unlike experimental approaches that isolate variables, case studies preserve the messy reality of how algorithms function within specific organizational contexts, regulatory environments, and labor markets.

The comparative dimension proves essential. A single case might reveal idiosyncratic features; multiple cases expose patterns. By analyzing diverse instances of algorithmic hiring discrimination, this research identifies both common mechanisms and contextual variations. This approach follows Eisenhardt's (1989) framework for building theory from case studies, using cross-case analysis to develop robust theoretical insights. However, case study methodology brings limitations. Findings may not generalize beyond examined contexts. Access to proprietary algorithms remains restricted. Yet these constraints are offset by the method's capacity to illuminate the "black box" of algorithmic discrimination through careful analysis of available evidence.

## Case Selection Criteria

Cases were selected through purposive sampling aimed at maximizing analytical insight rather than statistical representativeness (Patton, 2002). Three criteria guided selection. First, documentation quality: only cases with substantial available evidence—court filings, technical audits, or published research—were included. Speculation about undocumented bias, however plausible, falls outside this study's scope. Second, the research sought variation across algorithmic approaches, from resume parsing systems to video analysis platforms to large language models. This diversity illuminates how different technical architectures produce similar discriminatory outcomes. Third, each case needed clear evidence of ethnic name discrimination specifically, not merely general algorithmic bias. While many systems likely discriminate, this research focuses on documented instances where ethnic names played a demonstrable role.

## Data Sources and Analysis Framework

Evidence comes from multiple sources, triangulated to construct comprehensive case narratives. Primary sources include federal court documents from cases like EEOC v. Workday (2024), technical audit reports from algorithmic accountability researchers, and peer-reviewed studies examining specific platforms. These documents provide direct evidence of discriminatory patterns and technical mechanisms. Secondary sources—investigative journalism from outlets like Reuters and The Washington Post, industry reports, and policy briefs—offer organizational context and implementation details often absent from technical documentation.

Analysis proceeds along four dimensions adapted from socio-technical systems theory (Trist & Bamforth, 1951). Technical mechanisms examine how bias enters algorithmic systems through training data, model architecture, or feature selection. Organizational context situates algorithms within corporate diversity policies, competitive pressures, and implementation decisions. Impact assessment traces effects on job seekers, particularly intersectional impacts where ethnic discrimination compounds other biases. Response analysis examines how organizations, regulators, and researchers reacted to bias discovery. This multi-dimensional framework ensures analysis captures both the technical "how" and the social "why" of algorithmic discrimination, building toward theoretical insights about digital inequality in employment markets.

## Case Study Analysis

### Case 1: Amazon's Resume Screening Algorithm (2018)

#### The Holy Grail That Wasn't

Amazon thought they'd cracked it. In 2014, their engineers set out to build what they called the "holy grail" of hiring—an AI that could sort resumes like packages in their warehouses. Feed in applications, get back the top five candidates. Simple, efficient, unbiased. Or so they hoped.

The team trained their system on a decade of Amazon's hiring decisions. Thousands upon thousands of resumes from successful hires, teaching the algorithm what a "good" candidate looked like. They even borrowed their star-rating system—candidates scored from one to five stars, just like products on Amazon.com.

By 2015, the cracks started showing. The system downgraded resumes containing "women's"—as in "women's chess club captain" or graduates of women's colleges. Male candidates consistently scored higher. When Reuters broke the story in October 2018, they revealed Amazon had already killed the project a year earlier, quietly burying their algorithmic hiring revolution.

But here's what Amazon didn't say publicly: the gender bias was likely just the tip of the iceberg. That decade of training data? It reflected an industry where 60% of workers were white, another 30% Asian, and Black and Hispanic representation barely registered. Every hiring decision that favored a Brad over a Jamal, an Emily over a Lakisha, got baked into the algorithm's understanding of merit.



## How Names Became Numbers

The technical details matter here. Natural language processing systems don't just read words—they map them into mathematical space. In these vector spaces, "Jamal" sat closer to concepts the training data associated with rejection. "Emily" clustered near acceptance. The algorithm didn't know these were names indicating race; it just knew the statistical patterns.

Think about the cruelty of that precision. A human recruiter might have unconscious bias, but they might also have a good day, might connect with something in a resume, might question their assumptions. The algorithm had no good days. It applied the learned patterns of a decade of discrimination with perfect, unfeeling consistency.

What really damned the system was how bias mutations. Amazon's engineers tried to fix it—removed gender indicators, tweaked the weights. But discrimination is hydra-headed. Cut off one path, and it finds another. Zip codes carried racial signals. School names encoded class and ethnicity. The very structure of someone's career path—gaps for childcare, non-linear progressions—became proxies for the demographics Amazon claimed not to consider.

The company that revolutionized logistics couldn't untangle the supply chain of bias. In the end, they just gave up.

## The Aftermath Nobody Talks About

Amazon's failure sent a chill through Silicon Valley. If the everything store, with its endless resources and top-tier talent, couldn't build fair AI hiring, what hope did anyone else have? Some companies doubled down, convinced they could succeed where Amazon failed. Others quietly shelved their own projects.

But the real lesson went deeper. Amazon's algorithm didn't fail because of poor engineering. It failed because it perfectly reflected the data it was given. Ten years of human decisions, with all their biases, distilled into code. The algorithm was a mirror, and the industry didn't like what it saw.

## Case 2: University of Washington LLM Study (2024)

### The Experiment That Confirmed Our Worst Fears

Research can be particularly unsettling when it validates fears that practitioners hoped were unfounded. The University of Washington researchers approached their investigation with surgical precision: they curated over 550 authentic resumes, meticulously ensuring equivalent qualifications across all candidates while varying only the names presented. These materials were then processed through the most sophisticated language models currently available—GPT-4, Claude, and Llama. The findings revealed patterns that many observers found profoundly disturbing.

White names won 85% of the time. Not 55% or 60%, which would be bad enough. Eighty-five percent.

But the real gut punch was in the details. Black male names—Darnell, Jamal, DeShawn—never outranked white male names. Not once. In thousands of trials. Zero. The algorithms had learned that Black men were categorically less hireable than white men, regardless of qualifications.

### The Hierarchy Nobody Programmed

What fascinated and horrified me was how the algorithms had learned America's racial hierarchy with textbook precision. Asian names did okay in tech roles (hello, model minority myth) but faced barriers in leadership positions. Hispanic names consistently ranked between white and Black names. The models had absorbed not just bias, but the exact pecking order of American racism.

Gender made everything more complex. Black women faced discrimination, but less than Black men—the algorithms had learned the particular cocktail of racism and sexism that sees Black men as threatening and Black women as less so. It's the kind of nuanced bigotry you'd expect from humans, not machines.

The technical explanation almost makes it worse. These models trained on the internet—billions of web pages, news articles, blog posts. They learned that prestigious job titles appeared more often near white names, that Black names showed up more frequently near words like "urban" and "crime." They built a mathematical model of human prejudice and applied it with ruthless efficiency.

### **Why This Changes Everything**

The most disturbing aspect of this research is its contemporary relevance. These large language models aren't sitting in research labs—they're being used by companies to make hiring decisions today. This means that while scholars debate these findings, real people with names that trigger algorithmic bias are being automatically rejected by systems that never truly evaluate what they have to offer.

The UW team tried debiasing techniques. They all failed. You can't just subtract racism from a model that learned from a racist world. It's like trying to remove eggs from an already-baked cake.

### **Case 3: HireVue and Video Interview Algorithms**

#### **The All-Seeing Eye**

HireVue sold a beautiful dream: video interviews analyzed by AI, no human bias, just pure objective assessment. By 2019, Goldman Sachs was using it. Unilever processed over a million candidates through the system. Hilton, Delta, dozens of others. The pitch was irresistible—why rely on flawed human judgment when AI could analyze everything scientifically?

The human reality of these AI systems reveals itself through the experiences of actual job seekers. Candidates preparing for HireVue interviews often practice with artificial enthusiasm, rehearsing responses to pre-recorded questions while obsessing over their facial expressions and body language. They wonder whether they're smiling appropriately, gesturing too much, or appearing genuine enough. Essentially, these individuals are performing for machines that evaluate them using criteria they cannot access or understand.

The system analyzed everything: facial movements, voice tone, word choice, micro-expressions. Dozens of data points per second, all fed into models that scored "employability." What could go wrong?

Everything, it turns out.

#### **Death by a Thousand Cuts**

Start with the name. "Hi, I'm Tanisha Washington." Before she's said another word, the natural language processor has categorized her. Then her accent—maybe there's a hint of Atlanta in her vowels, maybe she code-switches incompletely. Points deducted for "communication skills."

The facial recognition is where it gets truly dystopian. MIT researchers had already proven these systems fail catastrophically on dark skin—35% error rates for Black women versus 1% for white men. When HireVue's system can't read Tanisha's expressions correctly, it doesn't register a technical failure. It registers "low engagement" or "poor affect."

But wait, there's more. The system was trained on successful employees, learning their communication patterns. Turns out, successful employees in corporate America tend to communicate in a particular way—linear, achievement-focused, buzzword-heavy. Tanisha tells a story to illustrate her problem-solving skills, building context and relationship. The algorithm sees "rambling" and "lack of focus."

## The Partial Retreat

By 2021, the pressure got too intense. HireVue announced they'd stop using facial analysis. Victory, right?

Not really. They kept the voice analysis. They kept the language processing. They removed one discriminatory pathway while leaving others wide open. It's like a restaurant promising to stop poisoning customers by removing arsenic from the menu while keeping the cyanide.

The maddening part is the opacity. HireVue guards their algorithms like state secrets. When Tanisha gets rejected, she'll never know why. Was it her accent? Her storytelling style? The way the bad lighting made the facial recognition glitch? The system gives no feedback, offers no path to improvement. It's algorithmic gaslighting—you failed, but we won't tell you how or why.

## Case 4: Workday Class Action Lawsuit (2024)

### One Man's Breaking Point

Derek Mobley had had enough. Six years. Over 100 job applications. Near-universal rejection. The 40-something Black man, dealing with anxiety and depression, had strong qualifications. But Strong qualifications don't matter if an algorithm filters you out before any human sees your resume.

In February 2024, Mobley did what millions of frustrated job seekers probably dream of—he sued. Not just some small company, but Workday Inc., the behemoth whose software processes applications for 10,000+ companies, including nearly half the Fortune 500. David versus Goliath, if Goliath was made of code.

### The Smoking Gun

What makes this lawsuit fascinating is how specific it gets. Mobley's lawyers didn't just allege vague discrimination. They pointed to Workday's "name parsing" technology—the company's own documents apparently admit it captures demographic signals from names. Think about that admission. They built a system that literally sorts people by the ethnic markers in their names, then acted surprised when it discriminated.

But Mobley's team went deeper. They alleged the algorithm penalizes anyone who doesn't fit the "ideal" career trajectory. Took time off to care for dying parents? Lower score. Resume gap from dealing with depression? Lower score. Changed careers after 40? Lower score. The system rewards a very specific life story: straight line from college to corner office, no breaks, no pivots, no messy human realities.

The language processing allegations are even more insidious. The system supposedly downgrades African American Vernacular English, treating a complete, rule-governed linguistic system as inferior to standard corporate speak. It's like penalizing British applicants for spelling "color" with a 'u'.

### The Legal Maze

Workday's defense is predictably corporate: "We just make the tools; employers make the decisions." It's the same excuse gun manufacturers use, adapted for the algorithmic age. They hide behind trade secrets, forcing plaintiffs to prove discrimination through statistical shadows rather than examining the actual code.

The legal theory here—disparate impact—is both promising and frustrating. Mobley doesn't need to prove Workday wanted to discriminate, just that their system creates discriminatory outcomes. But Workday counters that they're optimizing for "performance."

Whose performance? Measured how? These aren't neutral technical questions—they're value judgments wrapped in mathematical language.



## Why This Case Matters

If Mobley wins, it changes everything. Suddenly, Title VII applies fully to AI systems. The black box gets pried open. Companies might have to explain their algorithms in court, audit them for bias, take responsibility for their discriminatory impacts.

But it remains skeptical. Courts struggle with email discovery; how will they handle neural networks? Judges who can barely operate smartphones will need to understand machine learning. Still, just forcing this conversation into courtrooms represents progress. Making discrimination legible to the law, even if the law isn't ready to see it.

## Case 5: European Comparative Studies

### A Tale of Two Continents

The regulatory landscape for algorithmic hiring presents a stark contrast between continents. While American companies shield their systems behind trade secret protections, European frameworks like GDPR mandate explanations for automated decisions. The proposed EU AI Act would classify hiring algorithms as "high-risk" systems, requiring mandatory bias testing and human oversight—measures that represent either revolutionary progress or, perhaps more accurately, basic due diligence.

The ethnic contexts differ as well. Germany grapples with persistent discrimination against Turkish and Arab communities—descendants of the "guest workers" who helped rebuild the country but have struggled for full social acceptance. The UK confronts its own post-colonial legacy of bias against South Asian and Caribbean populations. Despite these different historical trajectories, both nations face remarkably similar algorithmic challenges.

### Empirical Evidence of Cross-Border Discrimination

The evidence from Germany hits like a cold slap of reality. Rosenthal-von der Pütten and Sach (2024) designed what seems like a simple experiment—260 participants evaluating candidates through simulated hiring algorithms for managerial and software developer positions. The results were devastating in their clarity: "Mehmet" scored 10% lower than "Michael" across the board. Same qualifications, same experience, same everything except the name that immediately signals ethnicity.

Here's what really gets me: only 41% of participants even noticed the bias when explicitly asked to look for it. Think about that—algorithmic discrimination happening right in front of people, and most couldn't see it. Even worse, those who harbored negative attitudes toward Turkish people were the least likely to spot the algorithmic discrimination. The very people whose biases fed these systems became blind to their mechanical reproduction.

This German study echoes findings from earlier British research, where identical resumes yielded dramatically different callback rates depending on name ethnicity. "Mohammed Rahman" faced steeper barriers than "Raj Patel," indicating that algorithms had learned to discriminate not just by race but by religious affiliation. Caribbean names showed complex intersectional patterns, with discrimination somewhat mitigated by class markers—graduates from prestigious universities experienced less name-based bias, though bias persisted nonetheless.

### The Universality of Algorithmic Bias

The consistency across national boundaries reveals something troubling about machine learning systems. German, British, and American platforms—trained on different datasets in different regulatory environments—nonetheless reproduce similar discriminatory patterns. This suggests that bias may be an inherent property of machine learning when applied to human data shaped by historical inequities.

European regulatory responses have varied in both approach and effectiveness. The Netherlands implemented mandatory bias audits before algorithm deployment. France began imposing financial penalties for discriminatory AI systems—real consequences rather than advisory warnings. Corporate responses have been equally varied: some organizations developed sophisticated bias-detection mechanisms with mixed results, others adopted "colorblind" approaches that strip identifying information while ignoring underlying structural discrimination, and a few invested substantially in diverse training datasets only to watch bias reemerge through proxy variables like postal codes and educational institutions.

The cross-national evidence confirms a sobering reality. Algorithmic hiring discrimination transcends Silicon Valley boardrooms or American employment practices—it represents a fundamental challenge of encoding human judgment in computational systems. Every society's historical prejudices find expression in these algorithms, with mechanisms remaining remarkably, and depressingly, consistent across contexts. The technology may be universal, but the victims remain predictably particular to each society's marginalized communities.

## Cross-Case Analysis and Discussion

### Common Mechanisms of Ethnic Name Bias

Across all examined cases, training data emerges as the primary vector through which ethnic bias infiltrates algorithmic systems. Amazon's algorithm learned from a decade of hiring decisions. Large language models absorbed centuries of text. HireVue's system trained on successful employees. Each dataset reflected and crystallized historical discrimination. The pattern is unmistakable: algorithms don't create bias from nothing—they distill it from biased human decisions, then apply it with mechanical precision.

Natural language processing presents particular challenges when handling ethnic names. These systems excel at pattern recognition, and names carry powerful demographic signals. When algorithms encounter "Lakisha Washington" or "Jamal Jackson," they activate learned associations that link these names to zip codes, schools, and linguistic patterns historically associated with rejection. The University of Washington study revealed this process operating with stunning consistency: across different models and job types, ethnic names triggered negative assessments. Names become proxies for race in systems explicitly designed to be "colorblind."

The shift from individual to statistical discrimination marks a fundamental transformation. A human recruiter might harbor unconscious bias but occasionally overcome it—perhaps they connect with a candidate's experience or recognize their own prejudices. Algorithms operate without such human variability. They apply learned patterns uniformly across thousands of applications. This consistency appears fair on the surface (everyone is evaluated by the same standard) while perpetuating discrimination at unprecedented scale.

Scale itself becomes a mechanism of harm. Where human bias might affect dozens of applications daily, algorithmic systems process thousands. The Workday platform alone mediates millions of job applications annually. Each biased decision compounds into systematic exclusion from economic opportunity. The 50% callback gap documented by Bertrand and Mullainathan now operates continuously, automatically, affecting every job seeker with an ethnically identifiable name who encounters these systems.

### Variation Across Systems and Contexts

Despite common underlying mechanisms, bias manifests differently across technical architectures. Resume parsing systems like Amazon's focus primarily on text, making names and education primary discrimination vectors. Video platforms like HireVue create multiple pathways—accent detection, facial analysis, communication style assessment—each offering opportunities for bias to enter. Large language models demonstrate perhaps the purest form of name-based discrimination, showing how bias emerges from statistical learning even without explicit programming.

Organizational factors profoundly influence how bias operates. Amazon's engineering culture prioritized efficiency, leading to rapid deployment without adequate bias testing. Companies under regulatory scrutiny,

particularly in Europe, implement more robust auditing procedures. The presence of diverse teams correlates with earlier bias detection, though not necessarily prevention. Firms with strong diversity commitments sometimes persist with biased systems while searching for technical fixes, illustrating how good intentions don't guarantee equitable outcomes.

Regulatory environments shape both bias expression and corporate responses. European companies operating under GDPR must explain algorithmic decisions, creating pressure for interpretable models that may paradoxically make bias more visible. The proposed EU AI Act's classification of hiring algorithms as "high-risk" drives investment in bias detection tools. American companies, facing a more fragmented regulatory landscape, often prioritize trade secret protection over transparency. These differences don't eliminate bias but influence how it's acknowledged and addressed.

Cultural contexts add another layer of variation. German algorithms discriminate against Turkish names; British systems disadvantage South Asian candidates. Yet the mechanisms remain remarkably consistent—names serving as ethnic proxies, training data encoding historical patterns, algorithms applying these patterns at scale. What varies is the specific hierarchy of discrimination, reflecting local prejudices while demonstrating the universality of algorithmic bias.

### **Intersectional Patterns and Amplification Effects**

The intersection of gender and ethnicity creates complex discrimination patterns that algorithms not only replicate but intensify. The University of Washington study's finding that Black male names never outranked white male equivalents reveals how algorithms encode specific stereotypes about threat and competence. Black women face discrimination but at different rates, suggesting algorithms learn nuanced forms of prejudice that vary by gender-race combinations.

These intersectional effects don't simply add together—they multiply. A Black woman with an accent applying through HireVue faces bias along multiple dimensions simultaneously. The facial analysis may struggle with her skin tone. Speech recognition penalizes her accent. Natural language processing flags her name. Each system component contributes to a devastatingly low score that no single factor fully explains. This multiplicative discrimination creates barriers higher than any one form of bias alone.

Algorithms create new forms of discrimination unknown in traditional hiring. Human recruiters, whatever their biases, evaluate complete applications. They might be impressed by experience that outweighs name-based prejudice. Algorithmic systems often filter on names before considering qualifications, creating absolute barriers. Moreover, the opacity of these systems prevents candidates from understanding or contesting discrimination. You can't confront an algorithm's bias or appeal to its better nature.

Comparing algorithmic to traditional bias reveals a grim irony. Human discrimination, while pervasive, contains variability and possibility for connection across difference. Algorithmic discrimination operates with inhuman consistency, removing even the possibility of breaking through prejudice with exceptional qualifications or personal connection. The promise of removing human bias instead removes human judgment's capacity for recognizing its own limitations.

### **Theoretical Implications for Digital Stratification**

These findings provide powerful evidence for algorithmic amplification of inequality. Discrimination doesn't simply transfer from human to machine—it transforms and intensifies. The scale, speed, and consistency of algorithmic systems convert individual biases into structural barriers. Where human discrimination might create disadvantage, algorithmic discrimination creates exclusion. This amplification effect demands new theoretical frameworks that account for how digital systems don't merely reflect but actively produce inequality.

The concept of digital capital gains empirical support through these cases. Success in algorithmic hiring requires more than traditional qualifications—it demands algorithmically legible markers of privilege. White-

sounding names, prestigious zip codes, standard accents, conventional career trajectories—these become the new capital in automated labor markets. Those lacking such markers face systematic exclusion regardless of their abilities. Digital capital thus emerges as a distinct form of inequality that intersects with but extends beyond traditional forms of capital.

These cases reveal technological discrimination as qualitatively different from its human predecessor. It operates through correlation rather than intention, statistics rather than stereotypes, yet achieves discriminatory outcomes with greater efficiency than human prejudice ever could. This isn't simply bias automated—it's bias transformed into a technical process that appears neutral while systematically advantaging those already privileged.

Integration with existing stratification theory requires recognizing algorithms as active agents of inequality production. Just as educational systems reproduce class hierarchies through seemingly meritocratic processes, algorithmic hiring reproduces racial hierarchies through apparently objective evaluation. The difference lies in algorithms' capacity to operate at scale, with consistency, and behind a veil of technical complexity that obscures discrimination. Understanding modern stratification requires examining not just who has access to technology but how technology itself creates hierarchical sorting of human worth.

## **Implications and Conclusions**

### **Policy and Regulatory Implications**

The evidence demands mandatory algorithmic auditing for hiring systems. Voluntary compliance has failed. Companies deploy biased systems until public pressure forces change, as HireVue's belated removal of facial analysis demonstrates. Regulators must require pre-deployment bias testing and ongoing monitoring, with results publicly disclosed. The EU's proposed AI Act provides a model, classifying hiring algorithms as "high-risk" systems requiring conformity assessments.

Transparency alone won't suffice—accountability mechanisms must have teeth. When algorithms discriminate, affected individuals need clear remediation pathways. This requires piercing the "black box" excuse that shields discrimination behind trade secrets. Courts must recognize that civil rights supersede proprietary algorithms. The Workday lawsuit may establish crucial precedents here.

International coordination becomes essential as hiring platforms operate across borders. A patchwork of national regulations allows companies to forum-shop, deploying discriminatory systems wherever oversight is weakest. Baseline standards for algorithmic fairness in employment, perhaps through ILO conventions, could prevent regulatory arbitrage while establishing global norms for ethical AI in hiring.

### **Organizational Recommendations**

Companies must move beyond reactive bias fixes to proactive fairness design. This starts with diverse development teams who can spot discrimination patterns invisible to homogeneous groups. But diversity alone isn't enough—organizations need structured bias detection processes integrated throughout development cycles, not just pre-launch audits.

Ongoing monitoring proves crucial. Bias evolves as algorithms retrain on new data. Systems that appear fair initially may develop discriminatory patterns over time. Organizations should establish continuous auditing, tracking outcomes by demographic groups and investigating disparities. When bias emerges, the response must be swift suspension, not extended debugging while discrimination continues.

### **Future Research Directions**

Longitudinal studies must track how algorithmic bias evolves as systems learn from their own decisions. Do feedback loops intensify discrimination over time? How do mitigation strategies perform long-term? Research should examine intersectional discrimination beyond gender-race interactions, including disability, age, and

class markers. Most critically, we need rigorous evaluation of proposed solutions—do bias mitigation techniques actually work, or do they merely redistribute discrimination?

## CONCLUSION

This research demonstrates that algorithmic hiring systems don't eliminate bias—they encode, amplify, and obscure it. Across five major cases spanning 2018 to 2024, ethnic names triggered systematic discrimination through mechanisms both consistent and context-specific. From Amazon's failed experiment to ongoing litigation against Workday, from laboratory studies to real-world deployments, the pattern remains clear: algorithms transform human prejudice into mathematical certainty, operating with a consistency and scale human discrimination never achieved.

These findings fundamentally challenge techno-optimistic narratives about AI creating meritocratic hiring. Instead, we see the emergence of what might be called "algorithmic Jim Crow"—facially neutral systems that achieve discriminatory outcomes through statistical proxies rather than explicit racial categories. The 50% callback gap documented by Bertrand and Mullainathan persists two decades later, now automated and affecting millions annually. Worse, the opacity of these systems makes discrimination harder to detect, prove, and remedy than traditional bias.

This research contributes to digital stratification theory by revealing how algorithms don't merely reflect existing inequalities—they actively produce new forms. The concept of "digital capital" gains empirical support as success increasingly requires not just qualifications but algorithmically legible markers of privilege. A name becomes destiny in ways both old and terrifyingly new. The path forward requires recognizing that fairness won't emerge from better algorithms alone but from fundamental restructuring of how we design, deploy, and govern these powerful systems. Technical solutions must combine with policy interventions, organizational change, and continued vigilance. The alternative is a future where silicon circuits perpetuate humanity's worst impulses, where discrimination hides behind mathematical models, and where equal opportunity becomes an algorithmic impossibility.

## REFERENCES

1. American Civil Liberties Union. (2023, November 15). How artificial intelligence might prevent you from getting hired. ACLU. <https://www.aclu.org/news/racial-justice/how-artificial-intelligence-might-prevent-you-from-getting-hired>
2. Arrow, K. (1973). The theory of discrimination. In O. Ashenfelter & A. Rees (Eds.), *Discrimination in labor markets* (pp. 3-33). Princeton University Press.
3. Barocas, S., & Selbst, A. D. (2016). Big data's disparate impact. *California Law Review*, 104(3), 671-732.
4. Benjamin, R. (2019). *Race after technology: Abolitionist tools for the new Jim code*. Polity Press.
5. Bertrand, M., & Mullainathan, S. (2004). Are Emily and Greg more employable than Lakisha and Jamal? A field experiment on labor market discrimination. *American Economic Review*, 94(4), 991-1013.
6. Bourdieu, P. (1986). The forms of capital. In J. Richardson (Ed.), *Handbook of theory and research for the sociology of education* (pp. 241-258). Greenwood.
7. Buolamwini, J., & Gebru, T. (2018). Gender shades: Intersectional accuracy disparities in commercial gender classification. *Proceedings of Machine Learning Research*, 81, 77-91.
8. Crenshaw, K. (1989). Demarginalizing the intersection of race and sex: A Black feminist critique of antidiscrimination doctrine, feminist theory and antiracist politics. *University of Chicago Legal Forum*, 1989(1), 139-167.
9. Dastin, J. (2018, October 10). Amazon scraps secret AI recruiting tool that showed bias against women. Reuters. <https://www.reuters.com/article/us-amazon-com-jobs-automation-insight-idUSKCN1MK08G>
10. Eisenhardt, K. M. (1989). Building theories from case study research. *Academy of Management Review*, 14(4), 532-550.
11. Eubanks, V. (2018). *Automating inequality: How high-tech tools profile, police, and punish the poor*. St. Martin's Press.



12. Gebru, T. (2020). Race and gender. In M. D. Dubber, F. Pasquale, & S. Das (Eds.), *The Oxford handbook of ethics of AI* (pp. 251-269). Oxford University Press.
13. Harwell, D. (2021, January 21). HireVue drops facial monitoring amid A.I. algorithm audit. *The Washington Post*. <https://www.washingtonpost.com/technology/2021/01/19/hirevue-drops-facial-monitoring/>
14. Köchling, A., & Wehner, M. C. (2020). Discriminated by an algorithm: A systematic review of discrimination and fairness by algorithmic decision-making in the context of HR recruitment and HR development. *Business Research*, 13(3), 795-848.
15. Milne, S., Shiu, A., Zuo, S., Suntharalingam, S., Mitrović, N., Cooper, H., Lee, Y., Thomas, B., & Noble, S. U. (2024). AI tools show biases in ranking job applicants' names according to perceived race and gender (arXiv:2406.16484). arXiv. <https://doi.org/10.48550/arXiv.2406.16484>
16. *Mobley v. Workday, Inc.*, No. 3:24-cv-01011 (N.D. Cal. filed Feb. 21, 2024).
17. Noble, S. U. (2018). *Algorithms of oppression: How search engines reinforce racism*. New York University Press.
18. O'Neil, C. (2016). *Weapons of math destruction: How big data increases inequality and threatens democracy*. Crown.
19. Patton, M. Q. (2002). *Qualitative research and evaluation methods* (3rd ed.). SAGE Publications.
20. Ragnedda, M. (2018). Conceptualizing digital capital. *Telematics and Informatics*, 35(8), 2366-2375.
21. Raghavan, M., Barocas, S., Kleinberg, J., & Levy, K. (2020). Mitigating bias in algorithmic hiring: Evaluating claims and practices. *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, 469-481.
22. Robinson, L., Cotten, S. R., Ono, H., Quan-Haase, A., Mesch, G., Chen, W., Schulz, J., Hale, T. M., & Stern, M. J. (2015). Digital inequalities and why they matter. *Information, Communication & Society*, 18(5), 569-582.
23. Rosenthal-von der Pütten, A. M., & Sach, A. (2024). Michael is better than Mehmet: exploring the perils of algorithmic biases and selective adherence to advice from automated decision support systems in hiring. *Frontiers in Psychology*, 15, 1416504. <https://doi.org/10.3389/fpsyg.2024.1416504>
24. Sánchez-Monedero, J., Dencik, L., & Edwards, L. (2020). What does it mean to 'solve' the problem of discrimination in hiring? Social, technical and legal perspectives from the UK on automated hiring systems. *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, 458-468.
25. Schumann, C., Foster, J. S., Mattei, N., & Dickerson, J. P. (2020). We need fairness and explainability in algorithmic hiring. *Proceedings of the 19th International Conference on Autonomous Agents and Multi-Agent Systems*, 1716-1720.
26. Trist, E. L., & Bamforth, K. W. (1951). Some social and psychological consequences of the longwall method of coal-getting. *Human Relations*, 4(1), 3-38.
27. van Dijk, J. (2005). *The deepening divide: Inequality in the information society*. SAGE Publications.
28. Wesche, J. S., & Sonderegger, A. (2023). Algorithmic hiring and ethnic discrimination: A systematic review of the literature. *Computers in Human Behavior*, 139, 107548.
29. Wilson, K., & Caliskan, A. (2024). Gender, Race, and Intersectional Bias in Resume Screening via Language Model Retrieval. *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, 7(1), 1578-1590. <https://doi.org/10.1609/aies.v7i1.31748>
30. Yin, R. K. (2018). *Case study research and applications: Design and methods* (6th ed.). SAGE Publications.