



EPIC-PP-2026-010 · RESEARCH ARTICLE · AI GOVERNANCE

The Algorithmic Wasta: AI Hiring and the Remaking of Social Capital in Gulf Labour Markets

EPIC – Emirati Perspectives and Insight Centre

Seeing what others miss...

20 October 2026

CC BY 4.0 · Creative Commons Attribution 4.0 International · creativecommons.org/licenses/by/4.0

The Algorithmic Wasta: AI Hiring and the Remaking of Social Capital in Gulf Labour Markets

EPIC – EMIRATI PERSPECTIVES AND INSIGHT CENTRE

EPIC-PP-2026-010 · AI Governance · Published 20 October 2026

© 2026 EPIC – Emirati Perspectives and Insight Centre. Released under Creative Commons Attribution 4.0 International (CC BY 4.0): <https://creativecommons.org/licenses/by/4.0/>

EPIC is an independent research initiative in formation. It is not yet a registered or licensed entity. EPIC does not accept donations, payments or contracts.

الواسطة الخوارزمية: التوظيف بالذكاء الاصطناعي وإعادة تشكيل رأس المال الاجتماعي في أسواق العمل الخليجية

Abstract

Objective: This article develops the concept of algorithmic wasta to explain how artificial intelligence (AI) hiring tools remake social capital in Gulf labour markets, where intercessory networks have long shaped access to jobs, promotion, and information. It asks whether automated screening weakens, absorbs, or disguises network advantage under labour-nationalisation regimes. **Methodology:** The study is a theoretical and integrative essay. It synthesises peer-reviewed research on wasta and social capital, empirical audits of bias in résumé screening and large language models, economic evidence on Saudi Nitaqat quotas, and selected comparative regulatory texts, including New York City Local Law 144 and the European Union AI Act. No new survey, interview, or administrative dataset is analysed, and claims about specific Gulf employers are avoided. **Results:** Four ideal-typical pathways are specified: displacement, in which transparent matching reduces brokers' informational advantage; translation, in which network signals such as referrals become algorithmic features; dual tracking, in which algorithms govern high-volume screening while networks allocate elite roles; and compliance wasta, in which networks mobilise to satisfy automated or quota metrics. **Conclusion:** AI hiring will not abolish social capital in the Gulf; it will remake it. The article recommends local fairness audits, bilingual Arabic–English validity testing, candidate notice and appeal, and nationalisation metrics that track career progression rather than headcount alone.

Keywords: *wasta; algorithmic hiring; social capital; Gulf labour markets; labour nationalisation*

الملخص

الأهداف: تطوّر هذه المقالة مفهوم «الواسطة الخوارزمية» لتفسير الكيفية التي تعيد بها أدوات التوظيف المعتمدة على الذكاء الاصطناعي تشكيل رأس المال الاجتماعي في أسواق العمل الخليجية، حيث أسهمت شبكات التوسط طويلاً في تحديد فرص الوصول إلى الوظائف والترقيات والمعلومات. وتساءل المقالة عما إذا كان الفرز الآلي يُضعف أفضلية الشبكات أو يستوعبها أو يخفيها في ظل سياسات توظيف الوظائف. **المنهجية:** الدراسة مقالة نظرية تكاملية، تجمع بين الأدبيات المحكّمة حول الواسطة ورأس المال الاجتماعي، والدراسات التجريبية التي رصدت التحيز في فرز السير الذاتية والنماذج اللغوية الكبيرة، والأدلة الاقتصادية حول برنامج نطاقات السعودي، إضافة إلى نصوص تنظيمية مقارنة مختارة، منها القانون المحلي رقم 144 لمدينة نيويورك وقانون الذكاء الاصطناعي الأوروبي. ولا تحلّل الدراسة بيانات ميدانية أو إدارية جديدة، وتتجنب إطلاق أحكام على جهات توظيف بعينها. **النتائج:** تحدد الدراسة أربعة مسارات نموذجية: الإزاحة، حين يقلّص الفرز الشفاف الأفضلية المعلوماتية للوسطاء؛ والترجمة، حين تتحول إشارات الشبكات مثل الترشيحات إلى سمات خوارزمية؛ والمسار المزدوج، حين تتولى الخوارزميات الفرز الواسع بينما تبقى الشبكات مهيمنة على المناصب العليا؛ وواسطة الامتثال، حين تتحرك الشبكات لاستيفاء المؤشرات الآلية أو حصص التوظيف. **الخلاصة:** لن يلغي التوظيف بالذكاء الاصطناعي رأس المال الاجتماعي في الخليج، بل سيعيد تشكيله. وتوصي الدراسة بمراجعات محلية للعدالة، واختبارات صلاحية ثنائية اللغة، وإشعار المتقدمين وحقهم في التظلم، ومؤشرات توظيف تتبّع التقدم المهني لا الأعداد وحدها.

الكلمات المفتاحية: الواسطة؛ التوظيف الخوارزمي؛ رأس المال الاجتماعي؛ أسواق العمل الخليجية؛ توظيف الوظائف

Introduction

In the political economy of Gulf labour markets, *wasta*—intercession through personal networks—has long structured access to jobs, promotions, and information. Parallel to that older technology of trust, firms across the region are adopting algorithmic tools for résumé screening, ranking, chat-based assessment, and workforce planning. This paper asks what happens when those two technologies of selection meet. It advances the concept of algorithmic *wasta*: the remaking of social capital through AI-mediated hiring, such that network advantages are translated into, displaced by, or disguised inside machine scores.

The argument is conceptual and comparative rather than based on a new survey. It synthesises three literatures that rarely speak to one another in Gulf studies: research on *wasta* as social capital; research on bias and governance in algorithmic hiring; and research on labour nationalisation regimes such as Saudi Nitaqat and UAE Emiratization. The Gulf is analytically productive for this encounter because formal institutions of localisation coexist with informal network practices, and because large expatriate workforces make nationality an unusually salient hiring variable—precisely the kind of variable that embedding models and ranking systems can learn as a proxy.

Three claims organise the analysis. First, *wasta* is not only corruption; it is a durable social-capital technology with bright and dark sides, as recent organisational research emphasises. Second, AI hiring tools do not simply 'replace favouritism with merit.' They re-

encode prior distributions of opportunity, language, and network signals unless deliberately constrained. Third, Gulf nationalisation policies create distinctive optimisation pressures: firms may use algorithms to satisfy quota metrics while informal networks continue to allocate the most valued roles, producing a dual track of formal algorithmic compliance and informal *wasta* allocation. The paper concludes with implications for researchers, employers, and regulators in Kuwait and the wider Arab Gulf.

Approach and Scope

This article is a theoretical and integrative essay. It does not report new interviews or administrative data. Its warrant is conceptual clarification under conditions of rapid technological change, when waiting for perfect datasets would cede the field to vendor white papers. Claims about specific Gulf employers' tools are therefore avoided unless supported by published research. Claims about *wasta* and about algorithmic bias are anchored in peer-reviewed studies and official regulatory texts cited below.

Scope conditions matter. The framework is developed for formal hiring in medium and large organisations—private firms, state-owned enterprises, and public agencies—where ATS and AI screening are plausible. It is not intended as a full theory of informal labour, domestic work, or day labour markets where *wasta* and apps take different forms. Within formal hiring, the framework is portable across GCC states but must be calibrated to each country's localisation law and public-sector size.

The analysis engages English-language and internationally indexed scholarship on *wasta* and algorithmic hiring. Arabic-language practitioner debates are acknowledged as important future sources; they are not systematically reviewed here. That limitation is stated so that readers in Kuwait do not mistake an English-indexed synthesis for an exhaustive map of local discourse.

Wasta, Social Capital, and Gulf Hiring

Defining Wasta without Caricature

A further clarification concerns scale. *Wasta* operates in micro-exchanges—a phone call to a hiring manager—and in macro-patterns that stratify who hears about vacancies at all. Ali and Weir's (2020) holistic model is valuable precisely because it refuses to reduce the phenomenon to interpersonal corruption alone. When AI systems ingest organisational histories, they ingest those macro-patterns as data. The ethical evaluation of the tool therefore cannot stop at the model's local accuracy; it must ask which social structure the model is amplifying.

Hutchings and Weir's (2006) comparison of *guanxi* and *wasta*, though older, remains useful for cautioning against both exoticisation and false equivalence. Network-based economic action is not unique to Arab societies; what varies is institutional channeling, moral discourse, and the presence or absence of countervailing formal HR systems. Gulf states' simultaneous investment in formal HR modernisation and persistence of relational hiring makes them a privileged site for watching conversion between social and algorithmic capital.

Cunningham and Sarayrah's (1993) classic study described *wasta* as a hidden force in Middle Eastern society, encompassing both intercessory mediation and the person who mediates. Later work has refined the concept. Barnett et al. (2013) analyse *wasta* through the economics of trust and cronyism under particular regulatory conditions. Tlaiss and Kauser (2011) show its importance for career success among Middle Eastern managers. Ali and Weir (2020) advance a holistic model bridging micro-level interpersonal exchange and macro-level institutional effects. Ali et al. (2024), drawing on interviews in Jordanian banking, show how trust and power position *wasta* inside recruitment and selection, with micro-level benefits for insiders and macro-level risks of exclusion and weakened formal institutions.

Al-Twal et al. (2024) summarise the dual character clearly: *wasta* can reduce transactional frictions and support loyalty and information flow, yet it can also distort hiring, promotion, and perceptions of fairness. Baranik et al. (2023), using survey evidence across Arab Middle Eastern contexts, link *wasta* to employment and income outcomes. The point for this paper is not to moralise *wasta* out of existence. It is to treat it as social capital in Bourdieu's (1986) sense—resources accessible through durable networks—and therefore as something that can be converted, misrecognised, or technologically transformed.

Social Capital Mechanisms Relevant to Hiring

Information asymmetry is the hinge. In Granovetter's classic account, weak ties matter because they carry non-redundant job information. Applicant-tracking systems and public job portals also carry information, but they carry it in formats that privilege certain literacy practices: keyword optimisation, platform accounts, and continuous online presence. Candidates embedded in professional WhatsApp groups or alumni lists may still learn of unadvertised roles first. Thus even when an algorithm ranks applicants for a posted job, *wasta* may have already shaped who applies. Algorithmic fairness metrics computed only on the applicant pool miss pre-pool exclusion.

Obligation and reciprocity further complicate automation. Coleman (1988) emphasises that social capital includes enforceable trust. A manager who hires a relative's neighbour may be discharging a debt that stabilises a wider network. An algorithm that blocks that

hire as 'nepotism risk' may be formally correct and socially disruptive, or it may be protecting merit. The evaluative standpoint matters. This paper does not defend obligation hiring; it insists that reformers who ignore obligation's social functions will design tools that are circumvented rather than obeyed.

Coleman's (1988) account of social capital emphasises obligations, information channels, and norms that facilitate action. Granovetter's (1973) strength-of-weak-ties argument shows how acquaintances can transmit job information that close circles do not. In Gulf labour markets, *wasta* often combines strong-tie obligations (family, tribe, school cohort, patron–client ties) with weak-tie brokerage (a well-placed intermediary who opens a door). Algorithmic hiring intervenes primarily at the information and screening stage: it ranks strangers at scale. That intervention can weaken the informational monopoly of brokers, or it can launder broker advantage into seemingly impersonal scores when CVs, referrals, and prior employers already encode network history.

Bourdieu's (1986) forms of capital help specify conversion. Social capital becomes economic capital when a tie yields a salary. Cultural capital—credentials, accent, university brand, polished CV English—becomes social capital when it unlocks networks, and becomes algorithmic capital when models reward those same signals. In Gulf private sectors dominated by expatriate professionals in many skilled roles, cultural capital often includes transnational education and English-medium corporate idioms. Nationalisation policies attempt to revalue national citizenship as a form of institutional capital. AI systems trained on historical hires may undervalue that institutional capital unless localisation criteria are explicitly introduced—and may overvalue legacy expatriate patterns if they are not.

Labour Nationalisation as Institutional Context

Nationalisation regimes also generate documentation practices—periodic reports, colour-band classifications in the Saudi Nitaqat system, and MoHRE declarations in the UAE—that are themselves digitised. Once digitised, they become trainable signals. A vendor can offer to predict which hires will improve a firm's band or rate. Prediction of regulatory metrics is not illicit; it becomes problematic when it encourages sham titles, rapid churn of citizen employees, or isolation of nationals in non-core units. Peck's (2017) evidence that quotas change behaviour should be read alongside qualitative work on how firms adapt strategically. Algorithms accelerate adaptation, including perverse adaptation.

Kuwait's labour market has its own localisation debates and public-sector employment patterns that differ from the Saudi and Emirati private-sector stories. Without claiming identical mechanisms, the algorithmic *wasta* framework still applies: wherever formal

localisation criteria meet informal mediation and new screening software, the four pathways can appear. Comparative Kuwaiti research should test which pathway dominates in government, oil-sector contractors, and SMEs respectively.

Gulf labour markets are shaped by nationalisation programmes that set quantitative and qualitative targets for citizen employment in the private sector. Peck's (2017) study of Saudi Arabia's Nitaqat programme provides causal evidence that hiring quotas can shift private-sector employment patterns, with complex effects on firm behaviour. The UAE's Emiratisation framework, supported by the Nafis programme and Ministry of Human Resources and Emiratisation (MoHRE) compliance tools, likewise uses targets, incentives, and penalties to raise citizen participation in private firms. Analogous localisation agendas exist in Kuwait, Qatar, Bahrain, and Oman, with country-specific designs and intensities.

These regimes matter for algorithmic hiring because they create measurable objectives that software vendors are eager to optimise: dashboards of nationalisation rates, predicted attrition of citizen hires, and 'skills gap' scores. Optimisation is not neutrality. A system that maximises a Nitaqat or Emiratisation metric may still allocate elite pathways through informal channels, or may hire nationals into roles with weak progression. Hertog's (2010) analysis of Saudi political economy reminds us that segmented labour markets and brokered state–business relations are sticky. Algorithms enter that stickiness; they do not automatically dissolve it.

Algorithmic Hiring: Promise, Bias, and Governance

What Hiring Algorithms Do

LLM-based evaluation introduces a new opacity. Unlike a logistic regression over structured features, an LLM may justify a low score with fluent prose that sounds reasoned while tracking spurious correlates—name ethnicity, formatting, or unpaid care-work gaps. An et al.'s (2025) controlled résumé study demonstrates that score gaps can translate into hiring-probability gaps at decision thresholds. Wilson and Caliskan (2024) show that retrieval-based screening can embed intersectional harms. For bilingual Gulf CVs, an additional hazard is English-centric tokenisation and training: Arabic content may be underweighted or inconsistently parsed, disadvantaging candidates who present strengths in Arabic.

Game-based and video assessments add biometric and behavioural data trails. Even where vendors claim proprietary validity, Selbst et al.'s (2019) critique of fairness abstraction applies: a game that predicts 'grit' in a US college sample may not predict per-

formance in a Kuwaiti customer-service role. Without local criterion validation, such tools risk becoming expensive randomness with a scientific aesthetic.

Contemporary hiring AI typically performs one or more of the following: parsing and embedding résumés; matching candidates to job descriptions; ranking or scoring; chatbot screening; game- or video-based assessment; and, increasingly, large language model (LLM) evaluation of cover letters and interview transcripts. Vendors promise speed, consistency, and relief from overloaded human screeners. In a Saudi labour-market context, Aleisa et al. (2023) propose an AI recruiting model (AIRM) architecture using clustering, sentence transformers, and ranking, reporting substantially faster shortlisting than human experts in their evaluation setting. Such systems address real frictions in high-volume markets.

Yet the fairness literature cautions that speed is not the same as justice. Raghavan et al. (2020) examine vendor claims about mitigating bias in algorithmic hiring and find uneven validation practices and limited transparency. Amazon's widely reported abandonment (Reuters, 2018) of an internal résumé-screening tool that penalised signals associated with women remains a cautionary case about training on historically skewed hires. An et al. (2025) show that several major LLMs exhibit gender and racial biases in automated résumé evaluation, with measurable effects on implied hiring probabilities. Wilson and Caliskan (2024) document gender, race, and intersectional bias in résumé screening via language-model retrieval. The lesson is not that every Gulf deployment is similarly biased—few public audits exist—but that bias is a demonstrated property of the tool class, not a speculative risk.

Sociotechnical Failure Modes

A third failure mode is accountability laundering. When a rejected candidate asks why, the recruiter points to 'the system'; the vendor points to the client's configuration; the client points to nationalisation or budget rules. Kellogg et al.'s (2020) contested-terrain framing predicts such displacement of responsibility. Personal-data laws that grant access and contestation rights can force a documented chain of responsibility; soft ethics pledges cannot.

Fourcade and Healy's (2017) insight that classifications allocate life chances is especially sharp where visa status and citizenship already allocate rights. An automated ranking that sorts 'sponsorship cost' or 'likelihood to relocate' may appear economically rational while reproducing structural inequality between citizens, long-term residents, and temporary workers. Whether such features are lawful or desirable is a political question; the algorithmic *wasta* lens insists it be asked openly rather than buried in a feature weight.

Selbst et al. (2019) warn against fairness abstraction: metrics that ignore social context produce solutions that fail when re-embedded in institutions. Kellogg et al. (2020) analyse algorithms at work as a contested terrain of control, where measurement reorganises power. Fourcade and Healy (2017) describe how markets see through classifications that allocate life chances. Translated to Gulf hiring, an embedding model 'sees' nationality, university, and prior employer as geometry. If historical data over-represent particular nationalities in engineering roles, the geometry will recommend more of the same unless constraints intervene.

A second failure mode is metric fixation under nationalisation. If compliance officers are rewarded for headcount of citizen hires, vendors will optimise headcount. Progression, learning, and retention may lag. Wasta then re-enters as the informal system that decides which nationals receive mentorship and which expatriates retain control of technical gatekeeping. The algorithm appears modern; the career tournament remains networked.

Emerging Regulatory Templates outside the Gulf

Ajunwa's (2023) broader analysis of quantified workers underscores that workplace surveillance and ranking technologies redistribute power under the banner of efficiency. Although her primary empirical base is not the Gulf, the legal-theoretical caution travels: consent to being scored is rarely meaningful when scoring gates livelihood. Candidate 'consent' checkboxes in ATS portals should not be mistaken for substantive authorisation, especially for migrant workers whose relationship to employers is already asymmetric.

Comparative regulation offers imperfect but useful templates. New York City's Local Law 144, with enforcement beginning 5 July 2023 (New York City Department of Consumer and Worker Protection [DCWP], 2023), requires independent bias audits and candidate notice for automated employment decision tools used in covered hiring and promotion. The European Union's AI Act (Regulation (EU) 2024/1689; European Union, 2024) classifies AI systems used in recruitment and worker management as high-risk under Annex III, triggering conformity duties. Gulf regulators have not, as of the sources consulted for this study, enacted equivalent binding AI-hiring audit statutes of general application. Soft-law AI ethics documents and personal-data laws provide partial coverage. That regulatory gap is an opportunity for Kuwait and peers to design context-sensitive rules before vendor practices harden.

The Algorithmic Wasta: A Conceptual Framework

Definition

Algorithmic wasta names the set of processes through which AI-mediated hiring remakes social capital in settings where intercessory networks already structure opportunity. It is not a claim that machines have kinship. It is a claim that machine selection interacts with network selection along identifiable pathways. Four ideal-typical pathways are distinguished below. Real organisations may combine several.

Pathway 1: Displacement

Displacement is most plausible in high-volume, standardised roles with clear credential requirements and strong public posting norms. It is least plausible where trust, security clearance, or family reputation are intrinsic to the role as currently institutionalised. Policymakers who announce that 'AI will end wasta' without specifying role types are selling a pathway as a universal law.

In displacement, transparent, well-governed matching systems reduce the informational scarcity that made brokers valuable. A graduate without a powerful uncle but with verified skills becomes visible in a ranked pool. Aleisa et al.'s (2023) AIRM vision gestures at this possibility for the Saudi market: faster, broader matching. Displacement is the pathway nationalisation reformers and anti-corruption campaigns hope for. It requires more than software: verified credentials, appeal rights, and human review that does not reintroduce unchecked vetoes by connected managers.

Pathway 2: Translation

Referral boosts illustrate translation cleanly. Many global ATS products treat employee referrals as positive signals because referred workers sometimes show better retention in vendor datasets. In a wasta-dense environment, referral is precisely the formalisation of intercession. Encoding it as a virtue score converts social capital into algorithmic advantage while waving the banner of evidence-based HR. Organisations serious about displacement should cap or audit referral boosts, not celebrate them uncritically.

In translation, wasta inputs are converted into algorithmic features. A referred candidate is tagged 'employee referral,' a feature that many applicant-tracking systems already boost. A prestigious local school or a known family business appears in the employment history embedding near successful past hires. The intermediary's work is laundered into

a score. Observers see meritocratic ranking; the underlying capital remains social. Translation is especially likely where vendors market 'culture fit' predictors trained on incumbent employees—incumbents who were themselves often hired through networks.

Pathway 3: Dual Tracking

Dual tracking may be stable because it serves powerful interests. Executives retain discretion over strategic hires; HR reports automation KPIs; ministries see rising platform usage. The equilibrium is politically comfortable and analytically deceptive. Research designs that sample only automated pipelines will overstate meritocratisation.

In dual tracking, algorithms govern high-volume, lower-stakes screening while wasta continues to allocate scarce, high-trust roles. Call-centre and entry graduate pipelines become automated; leadership programmes and sponsor-sensitive positions remain relationship-driven. Dual tracking can look like modernisation in HR dashboards while leaving the commanding heights of the firm untouched. For citizen jobseekers, dual tracking may mean easier access to first jobs and continued blockage at the glass network ceiling.

Pathway 4: Compliance Wasta

Coaching markets are an observable symptom of compliance wasta. Where assessment bots spread, so do tutors who sell keyword packs, interview-script rehearsal for video AI, and 'nationalisation-ready' CV templates. Some coaching is legitimate skill-building; some is adversarial evasion. Regulators should watch for proxy discrimination against candidates who cannot afford coaching, a new class cleavage inside ostensibly open portals.

In compliance wasta, networks mobilise to satisfy algorithmic or regulatory metrics. A firm under localisation pressure may use AI to identify citizen candidates who meet minimum thresholds, then rely on wasta to place relatives or clients into those counted roles. Alternatively, well-connected intermediaries may coach candidates to game keyword parsers and assessment bots. Here the algorithm does not eliminate wasta; it becomes a new surface on which wasta operates—much as bureaucratic forms once became surfaces for intercession.

Why the Gulf Intensifies These Pathways

A fifth intensifier is platform concentration. If a small number of government-linked portals intermediate large shares of applications, their ranking defaults become de facto industrial policy. That concentration is an opportunity for public-interest design—open au-

dit logs, dialect-aware Arabic NLP, disability accessibility—and a risk if procurement favours opacity. Social scientists should treat portal governance as public administration research, not only as HR tech.

Several structural features intensify algorithmic *wasta* in Gulf settings. First, nationality is a legally and politically salient variable, not a hidden correlate. Models can learn it from names, passport fields, visa status, and education trajectories. Second, labour markets are highly transnational; training data from global vendors may embed OECD career patterns poorly aligned with Gulf credentialing and Arabic–English bilingual careers. Third, personal-data and AI-governance regimes are still consolidating, so audit obligations are thinner than in NYC or the EU. Fourth, the moral evaluation of *wasta* is contested: what outsiders call nepotism, insiders may call obligation and trust repair in low-trust formal systems (Barnett et al., 2013). Tools that ignore that contestation will be gamed or rejected.

Remaking Social Capital: Winners, Losers, and Ambiguous Cases

Potential Winners

Another potential winner is the organisation that uses algorithms as decision support rather than decision replacement: shortlists generated by models, final choices by diverse panels with recorded reasons. That hybrid does not eliminate bias, but it creates artefacts for appeal and research. The operational point is a documented second look.

Jobseekers rich in verified skills but poor in brokers stand to gain under displacement—if displacement is real. Citizens targeted by nationalisation incentives may gain visibility when localisation constraints are hard-coded into ranking. HR departments drowning in volume may gain time. Vendors and consultants gain a market.

Potential Losers

Candidates whose excellence is relational, dialectal, or portfolio-based rather than CV-keyword-based may lose under brittle parsers—an issue acute for Arabic résumé formats and mixed-language documents. Groups historically underrepresented in training data—whether by gender, nationality, or disability—face the risks documented in the bias literature. Traditional brokers lose informational rents under displacement but may gain new rents under translation and compliance *wasta* by specialising in 'AI coaching.'

Ambiguous Cases: Women and Nationalisation

Nationalisation can also interact with gender when citizen women enter private sectors previously coded as expatriate-male. Algorithms trained on old data may underrate them; nationalisation constraints may force their inclusion at screening while translation pathways stall their promotion. Longitudinal data on promotion after algorithmic hiring adoption would be especially valuable for Kuwaiti and UAE labour research.

Tlaiss and Kauser's (2011) work on *wasta* and career success intersects with gendered network access. If men's networks have historically been denser in certain sectors, displacement could help women; translation of incumbent patterns could harm them. LLM résumé studies outside the Gulf show complex gender effects that are not uniformly anti-women, reminding us that bias direction is empirical, not a priori. Gulf-specific intersectional audits—by gender, nationality, and disability—are largely still to be done. Until they are, strong claims about automated equity should be treated as marketing, not evidence.

Implications for Kuwait and Comparative Gulf Policy

Research Implications

Methodologically, correspondence audit studies must be designed with ethics-board oversight and without imposing undue burdens on real employers. Synthetic application packets submitted to research sandboxes operated with employer consent may be preferable to deceptive mass applications. Where deception is used, debriefing and minimisation standards from experimental sociology should apply. The absence of such studies in the Gulf is itself a finding about research infrastructure.

Kuwaiti and regional social scientists can treat algorithmic hiring as a site for mixed-methods research: audit studies with synthetic CVs varying nationality and gender markers; ethnographies of HR teams adopting vendor tools; and network surveys that measure whether broker value declines after ATS adoption. The Journal of the Social Sciences' interest in contemporary Arab social challenges makes this agenda a natural fit. Conceptual work on algorithmic *wasta* should be stress-tested against Kuwaiti private-sector and public-sector hiring separately, given their different nationalisation and sponsorship dynamics.

Employer Practice

Public employers and state-owned enterprises should go further: publish aggregate pass-through statistics by gender and nationality for automated stages; procure systems that allow onshore processing of candidate personal data consistent with applicable privacy law; and prohibit vendor clauses that forbid research audits. If a tool cannot be audited, it should not rank citizens' life chances.

Employers that wish to avoid compliance wasta theatre should: (a) document the business necessity of each automated feature; (b) monitor adverse impact by nationality and gender on screening pass-through rates; (c) keep meaningful human review with written reasons for override; (d) separate localisation dashboards from promotion and mentorship accountability so that citizen hires are not dead-ended; and (e) test Arabic and bilingual document handling explicitly. Vendor claims of 'bias-free' models should be treated as hypotheses requiring local validation (Raghavan et al., 2020).

Regulatory Options

Regional cooperation could produce a GCC model code for automated hiring in government contractors—lighter than the EU AI Act, stronger than silence. Shared Arabic fairness benchmarks and a roster of independent auditors would lower costs for smaller states such as Bahrain and Kuwait. Policy think tanks and university journals both have roles: one for policy design, one for empirical validation.

Kuwait need not copy NYC Local Law 144 or the EU AI Act wholesale. It can, however, borrow proportionate ideas: notice to candidates when automated ranking materially affects them; periodic independent audits for large employers and government contractors; preservation of appeal channels; and restrictions on opaque 'culture fit' scores lacking construct validity. Personal-data rules should clarify lawful bases for processing biometric and behavioural assessment data. Nationalisation ministries should scrutinise whether AI-optimised compliance improves career quality, not only counts.

What Not to Do

Three errors recur in public debate. The first is romanticising wasta as harmless culture while ignoring exclusion. The second is romanticising AI as pure merit while ignoring training-data politics. The third is assuming that nationalisation metrics plus software equal inclusion. Algorithmic wasta analysis rejects all three. It asks, instead, which pathway—displacement, translation, dual tracking, or compliance wasta—a given deployment is actually producing.

An Illustrative Scenario (Not an Empirical Case)

Consider a stylised mid-sized Gulf services firm adopting an LLM résumé ranker while under pressure to raise citizen employment. The vendor configures a localisation boost and a referral boost. Citizen applicants without referral tags rise modestly in the ranking; referred expatriates with dense keyword overlap still dominate the top decile for technical roles; citizen referrals fill many counted nationalisation slots in support functions. HR reports 'AI-enabled Emiratisation/Saudisation progress.' Managers continue to fill architect and client-lead roles through known networks. This scenario is not offered as a description of any named employer. It is a diagnostic story that combines translation, dual tracking, and compliance wasta—the mixture this paper predicts will be common unless deliberately disrupted.

The scenario also clarifies what evidence would falsify the framework. If, after adoption, referral effects shrink, citizen hires appear proportionally in high-progression technical grades, and candidates receive intelligible reasons for rejection, displacement would be empirically dominant. Gulf research should hunt for such positive cases as eagerly as for failures; the framework is a typology, not a counsel of despair.

Discussion: Merit, Trust, and the Politics of Legibility

Merit is often invoked as the opposite of wasta. The algorithmic wasta framework complicates that binary. Merit is itself institutionally produced: through schools that confer cultural capital, through English that functions as a corporate gatekeeping language, and through internships that are distributed via networks. When an algorithm rewards those signals, it is rewarding a sedimented history of advantage. Calling the output 'merit' without inspecting the sediment is ideological, not scientific.

Trust pulls the other direction. In settings where formal HR is distrusted—because of opacity, delay, or past favouritism—wasta can be experienced as a reliability technology. Barnett et al.'s (2013) economic reading of wasta under particular regulatory conditions captures part of this. An AI system that is equally opaque simply relocates distrust onto vendors and dashboards. Legibility is therefore a design requirement: candidates and overseers must be able to understand, in plain language, which features mattered. Local Law 144's notice and audit requirements are one template for manufacturing legibility; Gulf adaptations should be bilingual and role-sensitive.

There is also a politics of who gets to define the target variable. If the model predicts 'successful employee' defined as twelve-month retention, it will favour risk-averse profiles and may punish women who exit for care reasons or citizens who rotate through nationalisation programmes. If it predicts supervisor rating, it will absorb supervisor bias. If it

predicts localisation contribution, it will optimise citizenship. Each target is a political choice dressed as engineering. Social scientists should insist that target choice appear in methods appendices and procurement files.

Finally, the framework invites humility about reform speed. Dual tracking and compliance wasta can absorb many generations of HR software. Displacement is possible but contingent. Celebrating portal launches as the end of favouritism repeats an old modernisation script. Measuring pathway shares over time is the more sober project.

Limitations

This article's limitations are explicit. It offers no new Gulf audit dataset; it under-represents Arabic-language scholarship; it treats GCC similarities at a level that blurs Kuwaiti particularities; and it cannot observe confidential vendor models. These limits bound the inference: the four pathways are sensitising concepts for research design, not estimated population parameters.

In practical terms, ministries that purchase or mandate AI screening for public employment should publish annual fairness summaries analogous to budget transparency reports. Without public numbers, algorithmic wasta remains a concept without a mirror. With public numbers, researchers and jobseekers can hold systems to account, and vendors face a market incentive to improve Arabic and localisation-aware validity rather than only English-centric benchmarks.

Conclusion

AI hiring will not abolish social capital in Gulf labour markets. It will remake it. Sometimes that remaking will widen opportunity by weakening informational gatekeeping. Sometimes it will launder old advantages into new scores. Sometimes it will automate the periphery while leaving the core to networks. Sometimes it will become a stage on which intercession performs compliance. Naming these pathways as algorithmic wasta gives researchers and policymakers a vocabulary that respects local social structure without surrendering to it, and that takes seriously the global evidence on hiring algorithms without pretending the Gulf is New York or Brussels.

The research programme ahead is empirical. The policy programme ahead is modest and urgent: transparency, local audits, bilingual validity, and nationalisation metrics that track careers rather than only headcounts. Social capital will remain. The question is whether Gulf societies will allow it to hide inside the machine.

References

- Aleisa, M. A., Beloff, N., & White, M. (2023). Implementing AIRM: A new AI recruiting model for the Saudi Arabia labour market. *Journal of Innovation and Entrepreneurship*, 12, Article 59. <https://doi.org/10.1186/s13731-023-00324-w>
- Ali, S., & Weir, D. (2020). Wasta: Advancing a holistic model to bridge the micro-macro divide. *Management and Organization Review*, 16(3), 657–685. <https://doi.org/10.1017/mor.2020.27>
- Ali, S., Weir, D., AlQahtani, M. S., & Mrabet, M. (2024). Emic and etic perspectives on HR practice for managing human resource issues affected by the prevalence of informal networks in Arab countries. *Administrative Sciences*, 14(10), Article 236. <https://doi.org/10.3390/admsci14100236>
- Al-Twal, A., Alawamleh, M., & Jarrar, D. M. (2024). An investigation of the role of Wasta social capital in enhancing employee loyalty and innovation in organizations. *Journal of Innovation and Entrepreneurship*, 13, Article 12. <https://doi.org/10.1186/s13731-024-00372-w>
- An, J., Huang, D., Lin, C., & Tai, M. (2025). Measuring gender and racial biases in large language models: Intersectional evidence from automated resume evaluation. *PNAS Nexus*, 4(3), Article pgaf089. <https://doi.org/10.1093/pnasnexus/pgaf089>
- Baranik, L., Gorman, B., & Wright, N. (2023). Wasta and its relationship to employment status and income in the Arab Middle East. *Current Sociology*, 71(5), 830–847. <https://doi.org/10.1177/00113921211024705>
- Barnett, A., Yandle, B., & Naufal, G. (2013). Regulation, trust, and cronyism in Middle Eastern societies: The simple economics of "wasta". *The Journal of Socio-Economics*, 44, 41–46. <https://doi.org/10.1016/j.socec.2013.02.004>
- Bourdieu, P. (1986). The forms of capital. In J. G. Richardson (Ed.), *Handbook of theory and research for the sociology of education* (pp. 241–258). Greenwood.
- Coleman, J. S. (1988). Social capital in the creation of human capital. *American Journal of Sociology*, 94(Supplement), S95–S120. <https://doi.org/10.1086/228943>
- Cunningham, R. B., & Sarayrah, Y. K. (1993). *Wasta: The hidden force in Middle Eastern society*. Praeger.
- European Union. (2024). Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). *Official Journal of the European Union*. <http://data.europa.eu/eli/reg/2024/1689/oj>
- Fourcade, M., & Healy, K. (2017). Seeing like a market. *Socio-Economic Review*, 15(1), 9–29. <https://doi.org/10.1093/ser/mww033>

- Granovetter, M. S. (1973). The strength of weak ties. *American Journal of Sociology*, 78(6), 1360–1380. <https://doi.org/10.1086/225469>
- Hertog, S. (2010). *Princes, brokers, and bureaucrats: Oil and the state in Saudi Arabia*. Cornell University Press. <https://doi.org/10.7591/9780801458774>
- Kellogg, K. C., Valentine, M. A., & Christin, A. (2020). Algorithms at work: The new contested terrain of control. *Academy of Management Annals*, 14(1), 366–410. <https://doi.org/10.5465/annals.2018.0174>
- New York City Department of Consumer and Worker Protection. (2023). Automated employment decision tools: Frequently asked questions. <https://www.nyc.gov/assets/dca/downloads/pdf/about/DCWP-AEDT-FAQ.pdf>
- Peck, J. R. (2017). Can hiring quotas work? The effect of the Nitaqat program on the Saudi private sector. *American Economic Journal: Economic Policy*, 9(2), 316–347. <https://doi.org/10.1257/pol.20150271>
- Raghavan, M., Barocas, S., Kleinberg, J., & Levy, K. (2020). Mitigating bias in algorithmic hiring: Evaluating claims and practices. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency* (pp. 469–481). ACM. <https://doi.org/10.1145/3351095.3372828>
- Reuters. (2018, October 10). Amazon scraps secret AI recruiting tool that showed bias against women. <https://www.reuters.com/article/world/amazon-scraps-secret-ai-recruiting-tool-that-showed-bias-against-women-idUSKCN1MK08J/>
- Selbst, A. D., Boyd, D., Friedler, S. A., Venkatasubramanian, S., & Vertesi, J. (2019). Fairness and abstraction in sociotechnical systems. In *Proceedings of the Conference on Fairness, Accountability, and Transparency* (pp. 59–68). ACM. <https://doi.org/10.1145/3287560.3287598>
- Tlaiss, H., & Kauser, S. (2011). The importance of wasta in the career success of Middle Eastern managers. *Journal of European Industrial Training*, 35(5), 467–486. <https://doi.org/10.1108/03090591111138026>
- Wilson, K., & Caliskan, A. (2024). Gender, race, and intersectional bias in resume screening via language model retrieval. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society* (pp. 1578–1590). <https://doi.org/10.1609/aies.v7i1.31748>
- Ajunwa, I. (2023). *The quantified worker: Law and technology in the modern workplace*. Cambridge University Press. <https://doi.org/10.1017/9781316888681>
- Hutchings, K., & Weir, D. (2006). Guanxi and Wasta: A comparison. *Thunderbird International Business Review*, 48(1), 141–156. <https://doi.org/10.1002/tie.20090>