The Courtship of Nations: AI Woos the Powerful While the Bill Goes Elsewhere

Weekly Analysis — https://ainews.social

There is a particular grammar that governs how artificial intelligence is sold to the people who can grant it entry — heads of state, general counsels, procurement officers, chiefs of police — and it is the grammar of seduction. The suitor arrives bearing superlatives. He promises transformation, a leap into "the age of AI," a competitive edge that cannot be refused. What he does not do, ever, is sign anything that would bind him should the transformation fail to arrive. This is the tell. A courtship is not a contract, and the week's social story is the story of courtship mistaken for, and marketed as, commitment. Watch the move closely and you notice that the promises flow upward, toward the powerful, while the obligations — the proof, the risk, the cost, the accountability — are quietly redistributed downward, toward the governed, the applicant, the arrested, the annotator earning less than two dollars an hour.

Consider the diplomatic weather. In the same season that the United States has decided to push what one survey of the regulatory landscape calls "a looser approach to AI regulation, while the EU pushes new law" [26], the industry has learned that the fastest way to secure a market is to charm the sovereign rather than negotiate with the citizen. The genius of seduction diplomacy is that it produces headlines about the future without producing clauses about the present. When a vendor promises to "push a country toward the age of AI," the promise is structurally unlike a bond: there is no maturity date, no default provision, no schedule of remedies if the jobs never materialize and the sovereignty never arrives. The nation keeps the infrastructure it was told it needed; the vendor keeps the revenue and the goodwill; and the gap between the two is filled with a language so warm it feels rude to ask for terms.

[26] US pushes looser approach to AI regulation, while EU pushes new law

This essay is about who does the asking, and who never gets to. Because if you map the AI discourse of a single week not by what it celebrates but by whom it addresses — who is being wooed, who is being managed, who is merely being processed — a stark asymmetry appears. The suitors court the powerful. The bill, reliably, goes elsewhere.

A Promise With No Default Clause

The purest specimen of the form is the high-level vendor visit: an executive lands, meets the president or the prime minister, gestures at a national destiny, and departs having committed to nothing measurable. Jobs are invoked but not guaranteed. Training is implied but not funded. Sovereignty over data and models is precisely the thing not on the table, because the entire business model depends on that sovereignty remaining with the vendor. What the executive offers is a mood; what he withholds is a mechanism. And the reason this works — the reason it keeps working — is that the interests of the courted official and the courting firm align beautifully at the level of the announcement and diverge only later, at the level of outcomes, when the official is out of office and the vendor is invoicing.

The structure repeats wherever AI is sold to elites as an efficiency dividend. When Wall Street banks began pressuring their outside law firms to cut fees on the theory that AI now does the associate's work, the story was framed as disruption's inevitable march [27]. But notice the direction of the pressure: it flows from the most powerful client toward the professional service provider, and from there — inevitably — toward the junior lawyers whose billable hours were the thing being compressed. The bank captures the savings. The technology captures the credit. The cost lands on the least powerful person in the transaction, who was never in the room where the promise was made. This is not a bug in the seduction; it is its purpose. The value flows to whoever holds the power to demand it, and the burden settles wherever there is the least capacity to refuse.

[27] Wall Street banks push Big Law to cut fees because of AI

Kate Crawford named the deep asymmetry here when she observed that "tech companies rarely suffer serious financial penalties when their AI systems violate the law and even fewer consequences when their ethical principles are violated" [7]. The courtship works precisely because there is no default clause — not in the diplomatic promise, and not, in any enforceable sense, in the law that is supposed to sit behind it. When one side can walk away from every commitment without penalty, the relationship was never symmetrical to begin with. It only looked like a negotiation. It was always a solicitation.

[7] The Atlas of AI - Power, Politics, and the Planetary Costs

And here is the divergence the regulatory map makes visible: while the American approach loosens, betting that seduction plus deregulation equals investment, the European approach at least attempts to write terms into law [26]. The distance between those two postures is the distance between courtship and contract. One trusts the suitor's good intentions; the other insists on getting them in writing. The vendor, unsurprisingly, prefers the courtship, and lobbies for the

[26] US pushes looser approach to AI regulation, while EU pushes new law

jurisdiction that offers it.

The Stakeholder Table Was Already Full

If seduction diplomacy is how AI is sold to nations, stakeholder convening is how it is sold to the conscience. The favored technology of inclusion in this discourse is the gathering — the policy hub, the advisory board, the multi-stakeholder consultation that assembles "experts from academia, government, and civil society" to deliberate the responsible path forward. These convenings are real, and they are earnest, and their exclusions are legible to anyone willing to read the guest list. Around the table where the governance of predictive policing is debated, there is one constituency reliably absent: the people who are actually policed by the systems in question. The subject of the algorithm is the missing stakeholder. The table was full without them.

This is not an accident of scheduling; it is the structural silence at the center of the whole enterprise. Consider what the record actually shows about how these systems land on the bodies of the governed. ProPublica's foundational investigation into criminal risk assessment demonstrated that the software used to predict recidivism was "biased against blacks," flagging Black defendants as future criminals at nearly twice the rate it did white defendants [20]. The people scored by that instrument were not consulted about its design, its thresholds, or its deployment. They discovered it the way one discovers weather: as a condition imposed from above, already in effect, non-negotiable.

[20] Machine Bias — ProPublica

The facial recognition record is more visceral still. The American Civil Liberties Union has documented "more than a dozen wrongful arrests due to police reliance on facial recognition technology," each one a person handcuffed on the say-so of a machine that misidentified them [30]. The pattern of these arrests is not that the technology failed in isolation; it is that officers treated a probabilistic match as probable cause, abandoning the corroboration standards they would have applied to any human tip. As reporting on these cases has shown, police "ignore standards after facial recognition matches," letting the machine's confidence substitute for the investigation it was supposed to assist [6]. The wrongfully arrested were, of course, never at the stakeholder table. They are the constituency whose exclusion the whole apparatus is designed, however unconsciously, to preserve.

[30] Wrongful Arrests Pile Up Due to Facial Recognition Technology

[6] Arrested by AI: Police Ignore Standards After Facial Recognition Matches

What makes the policing case the cleanest test of power in AI discourse is that the harm is documented with unusual precision and the remedy remains unusually absent. There is no shortage of analysis showing, as one account puts it, that AI "is automating injustice in American policing" and exposing "public safety as a sham" [4]. There is a comprehensive civil-liberties argument that "AI-powered surveillance is turning the United States into a digital police state" [5]. Even the open-source community has begun building tools like Fair-Code for "auditing algorithmic bias in criminal" justice systems [13]. The documentation is dense, rigorous, and growing. What it does not do — what it structurally cannot do from outside the room — is convert into a seat at the table where deployment decisions are made. This is the gap between harm documentation and solution-building, and it is not a gap of knowledge. It is a gap of power. We know exactly what is wrong. The people who know it best are the ones without standing to fix it.

[4] AI is automating injustice in American policing

[5] AI-Powered Surveillance Is Turning the United States into a Digital Police State

[13] GitHub - yakew7/Fair-Code

Crawford's history of how surveillance tooling migrates is the relevant frame here: the National Security Agency's "distinctive methods and tools have filtered down to classrooms, police stations, workplaces, and unemployment offices" [7]. The instruments of the powerful do not stay with the powerful; they trickle down onto the populations least able to contest them, arriving as ordinary administrative fact. Emory students who protested AI surveillance on their campus were reacting to precisely this migration — the arrival of a monitoring apparatus none of them had been consulted about [12]. Protest is what dissent looks like when consultation was never offered.

[7] The Atlas of AI - Power, Politics, and the Planetary Costs

[12] Emory Students Protest AI Surveillance on Campus

Relocating the Burden of Proof

The most consequential rhetorical move in the current workplace doctrine of AI is one almost nobody names out loud: the relocation of the burden of proof from the system to the person the system judges. In the old order, a hiring process that produced discriminatory outcomes bore the burden of justifying itself. In the emerging order, the algorithm is presumed neutral until an individual applicant can demonstrate — through litigation, through discovery, through expert statistical testimony most people cannot afford — that it erred. The default has flipped. The machine gets the benefit of the doubt; the human gets the homework.

The evidence that this flip is unjustified on the merits is now overwhelming. The largest study of AI hiring algorithms conducted to date found "clear racial disparities" in the outcomes these systems produce [17]. Stanford's Human-Centered AI institute has documented how "AI hiring tools can yield racial bias and systemic rejection," the word *systemic* doing crucial work — this is not a stray error but a

[17] Largest study of AI hiring algorithms to date finds 'clear racial disparities

structural tendency [3]. Reporting across jurisdictions has traced how ”automated hiring tools spark discrimination” complaints from candidates who never learn why they were filtered out [29]. The pattern is documented. What is contested is who has to prove it, and at whose expense.

[3] AI Hiring Tools Can Yield Racial Bias and Systemic Rejection

[29] Will AI give you the job? Automated hiring tools spark discrimination

The Workday litigation is the load-bearing case here. A court has allowed a discrimination claim to proceed against the vendor’s applicant-screening software, in effect telling millions of candidates that a piece of AI may have discriminated against them — and, crucially, placing on those candidates the task of establishing the fact [11]. The class action’s movement toward resolution has been tracked as a landmark [9], and legal commentators now speak openly of a rising wave — ”when machines discriminate,” the bias lawsuits multiply [28], so much so that one practitioner survey describes AI discrimination lawsuits as having ”exploded” [2]. But the litigation explosion is itself the symptom. Each lawsuit is a private individual, at private expense, attempting to do what the system was never required to do for itself: prove the algorithm wrong.

[11] Discrimination à l’embauche par l’IA : le procès Workday

[9] Class action lawsuit on AI-related discrimination reaches final

[28] When Machines Discriminate: The Rise of AI Bias Lawsuits

[2] AI Discrimination Lawsuits Explode: When Your Hiring Algorithm

And the law is not a clean instrument for the task. Legal scholarship has begun to map what it calls the antidiscrimination law dilemma — the awkward fit between statutes written for human discriminators and harms produced by statistical systems that discriminate without intent, through proxy variables and correlation rather than animus [25]. The dilemma is not incidental. It is the space into which the vendor’s presumption of neutrality expands. If the doctrine cannot easily name the harm, the harmed cannot easily win, and the burden stays where the courtship put it: on the person, never the product.

[25] Unpacking AI Bias and the Antidiscrimination Law Dilemma

This is where the workplace doctrine of ”foundational expertise” and the shift from assessing knowledge to assessing judgment reveals its underside. Framed as respect for human skill, it functions as a transfer of responsibility. The worker is told to develop the discernment to catch the machine’s errors — to become, in the corporate formulation, a thinking partner rather than a thinking substitute. But that formulation quietly makes the worker the error-correction layer for a system sold as labor-saving, and then holds the worker accountable when the correction fails. The professional must now prove the algorithm wrong on penalty of their own performance review. As one analysis of algorithmic management in delivery work has shown, the machine that assigns, scores, and disciplines does not deliberate with the worker; it degrades the worker’s health while presenting its judgments as optimization [19]. The partnership is rhetorical. The subordination is operational.

[19] Livreurs à domicile : comment le management algorithmique dégrade la santé des travailleurs

The same inversion is arriving inside the meeting room. When Microsoft's tooling can, as French reporting details, "surveil, profile, and judge" employees through the ordinary machinery of a video call — measuring participation, inferring engagement, generating assessments no one asked for — the worker is once again made the object of a system they did not consent to and cannot inspect [15]. The burden of proving the profile wrong, should it ever surface in a decision about their career, will fall on them.

[15] IA en réunion : quand Microsoft surveille, profile et juge vos salariés

The Bill Goes Elsewhere

Every seduction has an accounting, and the defining feature of the AI courtship is that the account is settled by people who were never party to the deal. The promise is made to the powerful; the invoice is delivered to the invisible. Nowhere is this clearer than in the human infrastructure the industry works hardest to keep offstage.

The systems that dazzle presidents are built on labor that would embarrass them. OpenAI's use of "Kenyan workers on less than two dollars per hour" to sift and label the traumatic content that makes a chatbot safe enough to demo is not an aberration of the model; it is a precondition of it [22]. The prouesses — the feats — of AI rest, as Swiss reporting puts it plainly, on "invisible workers exploited in poor countries" [10]. Journalism that has set out to uncover "the labor exploitation that powers AI" keeps arriving at the same structure: a value chain whose glamour concentrates at the top and whose degradation concentrates at the bottom, in the annotation farms and content-moderation queues of the global South [24].

[22] OpenAI Used Kenyan Workers on Less Than $2 Per Hour: Exclusive

[10] Derrière les prouesses de l'IA, l'exploitation de travailleurs invisibles

[24] Q&A: Uncovering the labor exploitation that powers AI

And the extraction is not static. The same generative systems built partly on Kenyan labor are now hollowing out the Kenyan digital economy that labor belonged to — "generative AI is making the economy of web workers wobble," as one account of the reversal describes it, automating away the very freelance and microtask income that the technology was assembled to provide [8]. Meanwhile the demand for cheap data labor migrates on — a report on why "Chinese AI seeks data in Kenya and who pays for it" shows the same country being mined again by a different suitor, the courtship of the powerful reproducing itself across the geopolitical field while the wage stays low and the seat at the table stays empty [31]. The question in that headline — *who pays for it* — is the question of the entire essay.

[8] Au Kenya, l'IA générative fait vaciller l'économie des travailleurs du Web

[31] ¿Por qué la IA china busca datos en Kenia y quién los paga?

The environmental account lands the same way. The carbon footprint of AI, examined against "figures, limits, and data quality," turns out to be systematically understated, its true cost obscured by the poor quality of the very data the industry chooses to release [14]. The pattern is by now familiar: the benefit is announced with precision, the cost is reported with fog. And the fog is not neutral. It falls on the communities nearest the data centers, the watersheds nearest the cooling, the grids nearest the load — populations who, once again, were not consulted and will not be compensated.

[14] Huella de carbono de la IA: cifras, límites y calidad del dato

There is a domestic version of this offloading, too, and it targets the people with the least margin to absorb it. When public assistance goes digital, the "hidden toll of the digital shift in social aid" turns out to be borne by claimants who face new delays, new errors, and new opacities in the systems meant to serve them [18]. Practitioners of social work debating whether AI can move them "toward a more" ethical practice are wrestling with a technology that arrives having already redefined the relationship between the institution and the person it processes [16]. In the welfare office as in the police station, the machine's efficiency is the institution's savings and the claimant's risk. The costs are individualized because individualization is the point.

[18] Le bilan caché du virage numérique de l'aide sociale

[16] Intelligence Artificielle et Travail Social : Vers une Pratique

Adoption Doing the Work Outcomes Used to Do

Here is the quiet scandal beneath the courtship, and it deserves to be stated without hedging: the celebrated transformation is, by the industry's own emerging metrics, failing to show up in the numbers, and the discourse has adapted by measuring adoption instead of outcomes. Diffusion is fast. Entry costs are low. Productivity gains — the whole ostensible reason for the seduction — remain stubbornly flat. When the promised payoff does not arrive, the industry does not revise the promise. It changes the subject to the impressive rate at which its product is spreading, as though the spread were itself the benefit. Adoption is now doing the argumentative work that outcomes used to do.

Once you see this substitution, the value flows resolve into sharp focus. The only reliably verified beneficiary of the courtship, so far, is the courtship itself — the vendor visit, the procurement contract, the fee the bank saves by squeezing the firm. When Wall Street pressures Big Law to cut fees "because of AI," the savings are real and immediate; whether the AI actually produced them, or merely furnished the pretext for a power play the banks wanted to make anyway, is a question the celebratory framing never asks [27]. The technology functions here less as a productivity engine than as a rationale — a story powerful enough to justify a transfer that was going to happen

[27] Wall Street banks push Big Law to cut fees because of AI

regardless.

The political system has begun, tentatively, to price in the possibility that the whole payoff fails. A proposed bill ”would tax AI tokens to fund jobs if the technology causes mass unemployment” — a striking piece of legislation precisely because it concedes, in its structure, that the promised prosperity may not arrive and someone will have to catch the fall [1]. But read it closely and even this contingency reveals the asymmetry. The tax is triggered *after* the mass unemployment, not before; it insures the harm rather than preventing it; it treats the destruction of livelihoods as a cost to be socially absorbed rather than a decision to be democratically made. It is, in its way, the most honest artifact of the week — a default clause written into law because everyone knows the courtship did not include one.

[1] A new bill would tax AI tokens to fund jobs if the technology causes mass unemployment

This is where the shape of the discourse itself becomes evidence. Survey the week's coverage and something close to two of every five substantive pieces are, at bottom, documentation of ethical failure — bias, wrongful arrest, exploited labor, surveillance overreach, welfare harm. That preponderance is not proof that AI is uniquely malignant; plenty of technologies generate plenty of failure. It is proof of where the *attention* has been forced to go. A field confident in its outcomes would be flooded with outcome data. A field whose outcomes are flat produces, instead, a flood of harm documentation and a counter-flood of seduction — the two dominant genres of the moment, and both of them symptoms of the same missing middle: verified, distributed, democratically-accountable benefit. Meredith Whittaker and Kate Crawford built the AI Now Institute on the premise that a ”more balanced view of AI is on the horizon,” one grounded in evidence rather than promise [7]. The balance has not yet arrived, but its absence is now loud enough to be a finding in its own right.

[7] Artificial Unintelligence - How Computers Misunderstand

Who Speaks, Who Pays, Who Escapes

The democratic stakes of this asymmetry are not abstract, and they are getting more concrete as the tools reach the ballot itself. Coverage of the 2026 municipal cycle warns of ”AI, deepfakes, and disinformation” straining democracy at exactly the moment citizens most need to trust what they see [21]. The same technology sold upward as a governance efficiency is deployed sideways as a persuasion weapon, and in both directions the person who bears the risk — the misinformed voter, the surveilled protester, the screened-out applicant — is the person whose consent was never sought. Even the housing market has required federal guidance on applying the Fair Housing Act

[21] Municipales 2026 : IA, deepfakes et désinformation, la démocratie

to advertising "through digital platforms," an acknowledgment that algorithmic targeting can reproduce the redlining the law was written to end, and that the harmed party must, once again, be defended after the fact rather than consulted in advance [23].

[23] PDF Guidance on Application of the Fair Housing Act to the Advertising of Housing Opportunities Through Digital Platforms

Return, finally, to the question of accountability — who gets blamed and who escapes. In the courtship model, blame follows the same downhill path as cost. When a hiring algorithm discriminates, the applicant must sue. When a facial recognition match is wrong, the arrested must clear their name. When the productivity gains fail, the worker must demonstrate their foundational expertise. When the data labor is exploited, the exploitation is offshored past the reach of the institutions that benefit. At each stage the entity that made the promise — the vendor, the bank, the department, the platform — recedes from view, protected by the fog it controls and the presumption of neutrality it engineered. Crawford's observation bears repeating because the whole architecture rests on it: firms "rarely suffer serious financial penalties when their AI systems violate the law and even fewer consequences when their ethical principles are violated" [7]. Escape is not a failure of the system. It is the system's most reliable feature.

[7] The Atlas of AI - Power, Politics, and the Planetary Costs

So who has power in AI discourse? The suitors — the vendors who set the terms, the executives who make the promises, the officials who are courted rather than the citizens who are governed. And who does not? The Kenyan annotator, the wrongfully arrested, the screened-out applicant, the welfare claimant, the delivery worker managed by a machine that will never explain itself, the voter targeted by a deepfake, the community absorbing the carbon and the water. These are the people the discourse addresses only as objects — processed, scored, exploited, offloaded — and never as parties. The social contract of AI is being drafted right now, in the language of seduction, and its striking feature is that it binds the governed while releasing the vendors. The promises float upward toward the powerful. The bill, as always, is mailed elsewhere — to the addresses of people who were never told a promise was being made in their name.

The remedy is not another convening, another hub, another table set with the same guest list. It is the insistence — legal, political, and moral — on the one clause the courtship has always refused to write: the default clause. The provision that names what is owed, and to whom, when the age of AI arrives for the powerful and passes everyone else by. Until that clause exists and binds, the courtship is not a partnership. It is a solicitation, conducted in the warmest possible language, at the expense of everyone not invited to reply.

---

1. [27] 2. [26] 3. [7] 4. [20] 5. [30] 6. [6] 7. [4] 8. [5] 9. [13] 10. [12] 11. [17] 12. [3] 13. [29] 14. [11] 15. [9] 16. [28] 17. [2] 18. [25] 19. [19] 20. [15] 21. [22] 22. [10] 23. [24] 24. [8] 25. [31] 26. [14] 27. [18] 28. [16] 29. [1] 30. [7] 31. [21] 32. [23]

[27] Wall Street banks push Big Law to cut fees because of AI
[26] US pushes looser approach to AI regulation, while EU pushes new law
[7] The Atlas of AI - Power, Politics, and the Planetary Costs
[20] Machine Bias — ProPublica
[30] Wrongful Arrests Pile Up Due to Facial Recognition Technology
[6] Arrested by AI: Police Ignore Standards After Facial Recognition Matches
[4] AI is automating injustice in American policing
[5] AI-Powered Surveillance Is Turning the United States into a Digital Police State
[13] GitHub - yakew7/Fair-Code
[12] Emory Students Protest AI Surveillance on Campus
[17] Largest study of AI hiring algorithms to date finds 'clear racial disparities
[3] AI Hiring Tools Can Yield Racial Bias and Systemic Rejection
[29] Will AI give you the job? Automated hiring tools spark discrimination

References

1. A new bill would tax AI tokens to fund jobs if the technology causes mass unemployment

2. AI Discrimination Lawsuits Explode: When Your Hiring Algorithm

3. AI Hiring Tools Can Yield Racial Bias and Systemic Rejection

4. AI is automating injustice in American policing

5. AI-Powered Surveillance Is Turning the United States into a Digital Police State

6. Arrested by AI: Police Ignore Standards After Facial Recognition Matches

7. Artificial Unintelligence - How Computers Misunderstand

8. Au Kenya, l'IA générative fait vaciller l'économie des travailleurs du Web

9. Class action lawsuit on AI-related discrimination reaches final

10. Derrière les prouesses de l'IA, l'exploitation de travailleurs invisibles

11. Discrimination à l'embauche par l'IA : le procès Workday

12. Emory Students Protest AI Surveillance on Campus

13. GitHub - yakew7/Fair-Code

14. Huella de carbono de la IA: cifras, límites y calidad del dato

15. IA en réunion : quand Microsoft surveille, profile et juge vos salariés

16. Intelligence Artificielle et Travail Social : Vers une Pratique

17. Largest study of AI hiring algorithms to date finds 'clear racial disparities

18. Le bilan caché du virage numérique de l'aide sociale

19. Livreurs à domicile : comment le management algorithmique dégrade la santé des travailleurs

20. Machine Bias — ProPublica

21. Municipales 2026 : IA, deepfakes et désinformation, la démocratie

22. OpenAI Used Kenyan Workers on Less Than $2 Per Hour: Exclusive

23. PDF Guidance on Application of the Fair Housing Act to the Advertising of Housing Opportunities Through Digital Platforms

24. Q&A: Uncovering the labor exploitation that powers AI

25. Unpacking AI Bias and the Antidiscrimination Law Dilemma

26. US pushes looser approach to AI regulation, while EU pushes new law

27. Wall Street banks push Big Law to cut fees because of AI

28. When Machines Discriminate: The Rise of AI Bias Lawsuits

29. Will AI give you the job? Automated hiring tools spark discrimination

30. Wrongful Arrests Pile Up Due to Facial Recognition Technology

31. ¿Por qué la IA china busca datos en Kenia y quién los paga?