

The Employment Game

Work as a socially embedded practice in the age of AI

Michael Stoyanovich

Version 1.8.1

Disclaimer

This paper is intended for informational and educational purposes only. The views and analyses presented - particularly those related to ethics, policy, and AI system design - reflect the author's interpretations and do not constitute legal, regulatory, or professional advice. Readers are encouraged to critically assess the content and consult appropriate experts or authorities before applying any concepts discussed herein. The author assumes no liability for any decisions or actions taken on the basis of this work.

AI may automate tasks, but it usually collides with the employment game.

When people ask how much of a profession an LLM can replace, they often begin with a sensible but incomplete move: break the profession into tasks, estimate which tasks a model can perform, and total up the percentage. The result sounds rigorous. It often is not.

The problem is not decomposition itself. The problem is what it leaves out. Much of this analysis treats a job as though it were exhausted by formal procedure, as if once the visible steps are mapped, the work has been adequately described. But that is precisely the picture of intelligence that earlier generations of AI got wrong. And it risks making the same mistake again - this time about work.

Jobs are not just bundles of tasks. They are socially, institutionally, and technically embedded practices. A role exists inside permissions, data dependencies, edge cases, escalation paths, compliance demands, accountability structures, and organizational norms. I call that larger structure the **employment game**.

I use the term *game* in the serious sense suggested by Wittgenstein's language games, not in the sense of playfulness or manipulation. The point is that neither language nor work is reducible to explicit rules alone. In both cases, formal instructions underdetermine practice. What matters is not only the rule, but its use: who may apply it, when it counts, what exceptions arise, how others interpret it, and what surrounding norms give it force.

There is also a coordination point here, closer to David Lewis. Jobs are not performed in isolation. They depend on shared expectations about who does what, when an output counts as final, when a case must be escalated, what evidence is sufficient, which exceptions are routine, and how others in the system are likely to respond. Work, in this sense, is not just rule-following. It is participation in a socially stabilized coordination structure. The employment game includes that common-knowledge layer.

Polanyi helps complete the picture. We know more than we can fully say, and much of work depends on that kind of practical knowledge: recognizing when a case is normal enough to process routinely, sensing when an exception matters, understanding how to phrase something so it will be accepted, or knowing what counts as sufficient evidence in a given institutional setting. That socially learned know-how is also part of the employment game.

That is why task automation and job replacement are not the same thing.

A model may produce an answer yet lack the authority to issue it. It may identify a pattern yet sit outside the liability structure required to act on it. It may draft a recommendation yet never enter the approvals, controls, and exceptions that make that recommendation institutionally valid. It may perform part of the visible work while remaining external to the employment game in which that work becomes real, accountable, and governable.

This is also where the “shock of the old” appears. In practice, the bottleneck is often not raw model capability. It is the inherited structure into which the model must be inserted: fragmented data, RBAC, workflow dependencies, regulatory intrusion, record-keeping, supervisory review, and organizational culture. The friction is not incidental. It is part of the employment game itself.

So the question is usually framed too narrowly. The real question is not simply whether a model can perform some percentage of tasks. The better question is whether it can participate in the employment game that makes those tasks meaningful, authorized, and accountable in the first place.

Working definition

Employment game: the coordinated technical, institutional, and social practice within which a job is actually performed.

It includes not only formal tasks, but also data dependencies, permissions, exception handling, handoffs, escalation paths, compliance requirements, accountability structures, practical know-how, and shared expectations about how actions are interpreted by others.

For that reason, task automation is not identical to job replacement.

Compact model

Task → Job → Employment Game

- **Task:** a bounded activity
- **Job:** a recognized bundle of responsibilities
- **Employment game:** the wider system of permissions, exceptions, coordination, accountability, and practical know-how within which the job is actually carried out

That yields the central claim: AI may automate tasks, may reshape jobs, but usually collides

with the employment game.

Five layers of the employment game

1. Formal layer

The written description of the work: procedures, SOPs, checklists, policy statements, job descriptions, KPIs.

2. Technical layer

The systems and data that enable or constrain action: source systems, integrations, workflow tools, RBAC, data quality, interfaces, logging.

3. Exception layer

The cases that break the script: ambiguity, incomplete information, overrides, appeals, special handling, judgment calls, local workarounds.

4. Institutional layer

The structures that make work legitimate and governable: regulation, auditability, retention, supervision, liability, sign-off, control requirements.

5. Social layer

The human setting in which the work is actually done: trust, role expectations, informal escalation, timing norms, status, institutional memory, culture, common knowledge about what counts as enough, final, appropriate, or review-worthy, and the know-how that rarely appears in formal documentation.

Thin automation analysis usually focuses on the first layer and glances at the second. But roles endure, and often resist substitution, because work is enacted across all five.

Two examples

Benefits administration

On paper, a benefits role may appear highly procedural. Determine eligibility. Review plan provisions. Process a transaction. Answer a participant question. Generate a notice. From a distance, this can look like ideal automation territory.

But the real work sits inside an employment game.

The formal layer includes plan rules, SOPs, templates, and required notices. The technical layer includes eligibility files, payroll feeds, historical records, document repositories, workflow systems, and tightly controlled permissions over protected data. The exception layer appears almost immediately: missing records, conflicting dates, unusual service histories, retroactive adjustments, plan-specific carveouts, and participant situations that do not fit the standard path. The institutional layer includes fiduciary obligations, retention expectations, privacy constraints where protected data is involved, and the need for auditable determinations. The social layer includes practical knowledge about how staff interpret ambiguous cases, when supervisors are consulted, how participant communications are phrased to avoid confusion, and what kinds of judgment are treated as routine versus

escalatory.

There is also a coordination structure at work. Staff know which cases can be handled locally, which require legal review, when an employer group must be consulted, and when a participant communication is considered safe to send. Much of this is not written as a rule for every case. It is maintained through shared expectations and common knowledge built over time.

An AI system may help summarize documents, draft participant communications, or identify likely next steps. It may reduce friction in the formal layer and parts of the technical layer. But that is not the same as replacing the role. The role is not just “apply rule to case.” It is participation in a larger practice of authorized, reviewable, and institutionally accountable administration.

That is the employment game.

Cybersecurity and GRC

A risk, compliance, or cybersecurity role can also look deceptively decomposable. Review evidence. Map controls. Draft risk statements. Assess gaps. Recommend remediation. Prepare a report. Again, the task inventory can make the work appear more substitutable than it really is.

But the actual role is embedded in an employment game.

The formal layer includes frameworks, control catalogs, policies, procedures, issue logs, and reporting templates. The technical layer includes systems inventories, evidence sources, ticketing platforms, IAM, scanning tools, configuration baselines, and document repositories. The exception layer includes compensating controls, incomplete evidence, inherited systems, conflicting stakeholder accounts, legacy environments, and situations where no framework requirement fits neatly without interpretation. The institutional layer includes regulatory expectations, contractual obligations, audit requirements, liability, review thresholds, and sign-off responsibilities. The social layer includes relationships with control owners, informal knowledge of which systems are genuinely risky, local credibility, escalation etiquette, and the practical difference between “documented compliant” and “operationally trusted.”

Again, coordination matters. Teams know which findings are serious enough to push upward immediately, which issues can be handled in the next review cycle, what level of evidence a particular auditor will accept, and how language in a report will be received by legal, operations, or executive leadership. These are not merely private intuitions. They are shared expectations that let the organization coordinate action.

An LLM may draft a control narrative, summarize evidence, or suggest likely risks. That can be useful. But the employment game is what turns a draft into an accepted control assessment, a flagged issue into an accountable remediation plan, or an observation into something leadership is willing to own. The model may support the work; it does not automatically inherit the authority structure, institutional legitimacy, or social trust that make the work count.

Task support is real. Job replacement is a different claim.

Why this matters

The phrase “AI can do 60 percent of the job” often mistakes the visible surface for the whole. A model may succeed at the formal layer while failing everywhere else that the job becomes institutionally consequential.

The more a role depends on exceptions, authority boundaries, regulated judgment, auditability, tacit coordination, institutional trust, and practical know-how, the more poorly it is described by task decomposition alone.

That does not mean AI has little impact. It means the impact is often misdescribed. AI can compress portions of work, alter role composition, shift where judgment is exercised, and increase the importance of review, exception handling, and governance. But those changes happen inside the employment game, not outside it.

None of this means the employment game is fixed forever. Institutions do adapt. Over time, organizations may redesign workflows, authority boundaries, review processes, and accountability structures to accommodate AI-in-the-loop as a normal part of work. But that is precisely the point: meaningful automation at the level of jobs usually requires not just better models, but changes to the employment game itself.

A practical diagnostic

To understand any role, ask:

1. What is the formal task?
2. What systems and data make it possible?
3. What exceptions routinely disrupt it?
4. Who is authorized to act, approve, or override?
5. What legal, regulatory, or audit burdens shape it?
6. What shared expectations and practical know-how make it actually work?

The answers will usually tell you more about automability than the task list alone.

Closing

Jobs are not just bundles of tasks. They are enacted within employment games. That is why replacing a task is not the same as replacing a job. And it is why the most consequential AI questions are rarely just about capability. They are about entry into systems of permission, coordination, exception, accountability, and practical knowledge that long predate the model itself.

Ethics, Disclosure, and Acknowledgements

Ethical Considerations

This essay does not draw on private, sensitive, or personally identifiable data. All examples are hypothetical, anonymized, or derived from public sources. No human-subjects research was conducted, and no institutional ethics review was required. The broader ethical implications concern public interpretation, policy design, and stakeholder responsibility in AI deployment. These implications are intended to provoke critical discussion and inform future regulatory and design frameworks.

Use of AI Tools

AI language models – most notably OpenAI’s ChatGPT – *were* used during the writing process as interlocutors: for brainstorming, structuring sections, and testing rhetorical clarity. These tools helped refine transitions, surface edge cases, and probe internal consistency. This meta-use aligns with the essay’s themes. Interacting with generative AI during authorship provided firsthand insight into the very limitations analyzed here—most notably fluency without grounding and responsiveness without responsibility at scale. Responsibility for all ideas, arguments, and conclusions lies solely with the human author.

Acknowledgements

Thank you to informal readers who offered critical feedback on earlier drafts. Their questions, challenges, and encouragement materially improved the final manuscript. Special thanks to those who pressed for clearer synthesis and for bridging philosophy and engineering as complementary perspectives on design. No institutional support, funding, or affiliation contributed to this work. All errors and omissions are the author’s alone.

Disclosure Statement

This work was conducted independently, without institutional affiliation, funding, or external influence. The views expressed are the author’s alone and do not represent any current or former employer. No financial or professional conflicts of interest are declared.

License & Attribution

This work is licensed under the Creative Commons Attribution 4.0 International (CC BY 4.0) license. You are free to share, adapt, and build upon this work for any purpose - including commercial use - so long as proper attribution is given. No additional permissions are required.

Full license terms: <https://creativecommons.org/licenses/by/4.0/>

Trademark Notice

The Four Philosophers Framework™ and The 4-Philosophers Framework™ are unregistered trademarks of Michael Stoyanovich. The CC BY 4.0 license does not apply to these trademarks. Use of the trademarked names is permitted for scholarly citation or descriptive reference but may not be used in connection with commercial products, services, or branding without permission.

How to Cite This Essay

Stoyanovich, Michael. *The Employment Game: Work as a Socially Embedded Practice in the Age of AI*. Version 1.8.1 (March 2026). <https://www.mstoyanovich.com>

Related Companion Papers

Stoyanovich, Michael. *Philosophy, Cognitive Science, and Policy: Interdisciplinary Perspectives on Generative AI from Wittgenstein, Lewis, Dennett, and Nagel*. Version 1.23.6 (December 2025). <https://www.mstoyanovich.com>

Stoyanovich, Michael. *The Human Lesson: A Response to Sutton through Wittgenstein, Lewis, Dennett, and Nagel*. Version 1.6.2 (November 2025). <https://www.mstoyanovich.com>

Stoyanovich, Michael. *The Question Concerning Learning: Babich, Heidegger, and the Enframing of Intelligence*. Version 1.0.1 (November 2025). <https://www.mstoyanovich.com>

Stoyanovich, Michael. *Context Collapse and the Four Philosophers: Wittgenstein, Lewis, Dennett, and Nagel in the Age of AI Chat*. Version 1.4.2 (November 2025). <https://www.mstoyanovich.com>

Stoyanovich, M. (January 2026). *A Constitution for a Chatbot: Reading Anthropic's "Claude's Constitution" through Context Collapse, the Four Philosophers, and the Shock of the Old* (Version 1.1.3). <https://www.mstoyanovich.com>

Version History and Document Status

This is a living document. As generative AI systems and their use evolve, this paper will be periodically updated to incorporate new empirical findings, theoretical insights, and policy developments. Major revisions are recorded here to preserve transparency and scholarly traceability.

Version	Date	Description
1.8.1	March 2026	Minor editorial update; added citation guidance and AI-in-the-loop adaptation paragraph.