Tomorrow's Affairs
Analysis of today
Assessment of tomorrow

By: Sami Mahroum

Should AI be taxed?

New York City Mayor-elect Zohran Mamdani's pledge to raise the minimum hourly wage to $30 captures a dilemma that is haunting most advanced economies. Even as wage floors rise, workers feel less secure.

Yet by responding with blunt tools like the minimum wage, policymakers are overlooking the deeper problem: the decoupling of human labor time from economic value.

For two centuries, labor markets were anchored in that relationship. Wages, contracts, and social protections all assumed that time was a reliable proxy for output. But AI has severed this link.

Medical diagnosticians who spent years mastering pattern recognition now compete with systems that process cases in seconds.

Lawyers with AI agents can complete tasks in minutes that would take hours otherwise.

The question is no longer whether minimum wages risk overshooting, but whether time-based compensation remains meaningful at all.

What is to be done? For starters, fiscal frameworks must recognize machine time as a distinct production input and price it accordingly.

This is not to suggest a "robot tax," but rather something closer to a metered input levy, using data that are already tracked in real time.

Robot tax

Trying to tax firms directly for their "degree of automation" has never worked, because such processes are organizationally opaque.

Algorithmic decision-making embedded in software is nearly invisible, and attempts to define a "robot" have proved legally absurd (does Excel count?).

Moreover, automation has long been deemed essential for maintaining global competitiveness, often posing a "problem" only for those offering manual labor.

> Policies that penalize productivity tend to be non-starters

Policies that penalize productivity tend to be non-starters. The most prominent effort came in 2017, when the European Parliament considered a proposal – embedded in MEP Mady Delvaux's report on robotics and AI – to levy a "robot tax" on companies that replaced workers with automated systems.

The idea was framed as a way to fund social protection and retraining, yet it was explicitly rejected in the final plenary vote amid concerns that the measure would undermine innovation and damage Europe's competitiveness.

Since then, robot-tax proposals have remained largely theoretical across EU member states, reinforcing the political reality that policies perceived as penalizing productivity struggle to advance.

AI-hours

The answer here may be to tax "AI-hours": the computational time consumed by AI systems performing economically productive tasks.

Since computational time is already precisely metered for billing purposes in the cloud industry, which underpins most enterprise AI services, AI-hours are one of those rare tax bases with built-in, automatically logged audit trails.

Consider a mid-size law firm in the United States that automates contract review with an AI system, replacing 25 full-time paralegals earning an average of $65,000 annually.

The hourly rate for this displaced labor is $32.50 (calculated as $65,000 divided by 2,000 standard work hours per year).

If the firm consumes 50,000 AI-hours annually in this automation, the taxable value of displacement is $1,625,000 (50,000 hours × $32.50).

A 15% levy on this displacement value would generate $243,750 in annual tax revenue.

> This framework taxes the economic value of human labor that AI systems replace

This framework taxes the economic value of human labor that AI systems replace, scales automatically based on wage levels across different occupations, and links the tax burden to the magnitude of labor displacement.

That may sound like a modest sum, but it would be meaningful when summed across all legal services, financial analysis, and medical diagnostics.

By 2028, enterprise AI infrastructure spending will reach $200 billion globally.

Augmentation and substitution

True, edge computing, which processes data near their source, would complicate things, because the usage logs are local.

But in these cases, the tax base would shift from data flows to proxies for capacity, such as chip specifications and energy consumption.

Since regulators already use these metrics for carbon-footprint reporting and energy audits, firms would simply need to register AI-capable hardware above defined thresholds and submit quarterly usage estimates, cross-checked against capacity and power-usage telemetry.

Of course, there is a risk that firms will game the system to qualify for lower tax rates. Clearly, definitions will matter enormously.

The two keywords requiring the most clarity are “augmentation” and “substitution.”

Augmentation refers to cases where humans use AI but remain the primary decision-making authority.

Substitution occurs when an AI executes full workflows, with humans intervening only under exceptional circumstances.

A radiologist reviewing AI-flagged anomalies but making final diagnoses has been augmented.

An AI system auto-approving 95% of loan applications is substituting for humans.

> With these terms clearly defined, an AI-hours levy could be tiered at, say, 5% for augmentation and 15% for substitution

With these terms clearly defined, an AI-hours levy could be tiered at, say, 5% for augmentation and 15% for substitution.

Workflow audits would determine classifications. If a firm claims augmentation but humans meaningfully handle fewer than 40% of cases, the system would be reclassified.

While the exact rates can be debated, the core principle mirrors the existing tax distinction between operating and capital expenditures.

And even with a 15% levy, AI-hours would still be dramatically cheaper than human labor.

Political backlash

The purpose of such a levy would be to nudge firms toward hybrid systems where human judgment improves outcomes, rather than toward fully autonomous configurations.

It would make labor displacement marginally costlier than augmentation without suppressing innovation. The industry itself should welcome such a policy.

The real threat to AI progress is not taxation

but political backlash. A predictable, rules-based levy could secure a social mandate for AI by ensuring that the gains are shared.

An AI-hours levy can convert machine productivity into public revenue without suppressing innovation

But what about the global dimension? As AI-hours become central to production, goods and services created with unpriced machine time could undercut those that are not.

Fortunately, we already have a way to deal with such arbitrage. A digital border adjustment mechanism (DBAM) could require firms to disclose AI-hour content for traded goods and cross-border digital services, mirroring carbon border adjustment mechanisms already in operation.

Implementation would not be difficult, because transfer-pricing rules already require multinationals to provide comparable documentation.

An AI-hours levy would have broad appeal. Businesses would get predictable rules instead of ad hoc wage politicking; labor would get a mechanism for capturing technological surplus (without blocking adoption); fiscal conservatives would get stability as payroll taxes erode; and progressives would get a tool to reduce inequality without having to resort to constitutionally dubious wealth taxes.

Coordination remains the main challenge. Early adopters may face short-term competitiveness pressures until DBAM regimes mature.

But plurilateral agreements – as with carbon pricing – could create a critical mass.

Debating whether minimum wages should be $16 or $30 misses the point. AI systems operate continuously at negligible cost.

Without a new fiscal framework, we will keep drifting toward an economy with a diminishing tax base.

An AI-hours levy can convert machine productivity into public revenue without suppressing innovation. It is the best response available.

Sami Mahroum, Founder of Spark X, previously held posts at INSEAD, the OECD, and Nesta.