AI Productivity Paradox, Fiscal Anchors, and the Measurement of Economic Resilience

A Systems Analysis of the Emerging Productivity Paradox

Bhargava R. Lakkaraju

Independent Researcher

Abstract

Artificial intelligence promises a generational leap in productivity, but the value created is not automatically captured by the economic actors who sustain demand. As long as human households anchor consumption, the true economic value of AI does not come from revenue expansion but from household cost-curve deflation. If the efficiency gains fail to translate into lower prices or improved financial resilience, the economy faces a structural paradox: productivity accelerates while purchasing power erodes.

This paper integrates IRS tax data, unemployment duration statistics, consumer credit segmentation, and sector-level income tiers to quantify how even a 1% downward mobility shock propagates through tax receipts, credit quality, household savings runway, and corporate exposure. We show that these effects are nonlinear and self-reinforcing, especially within prime-income groups that contribute disproportionately to federal revenues and the purchasing power.

To measure household resilience, we develop three new indicators: the Runway-to-Reemployment Ratio (R2R), the Equal-or-Superior Income Landing Rate (ESIL), and the Wage Retention Ratio (WRR). These metrics quantify the extent to which displaced workers can maintain consumption, avoid downward mobility, and recover earning power. Results show that for middle- and upper-middle-income groups—the stabilizing core of the fiscal base—R2R is already at or below parity under stress scenarios.

The analysis concludes that AI-first companies capable of passing cost efficiencies directly to consumers—especially in high-cost legacy sectors such as financial services, healthcare, airlines, and insurance—will not only create sustainable growth but also counteract the downward-mobility dynamics outlined here. In the absence of policy intervention, firms that reroute AI-driven efficiency gains into substantial price reductions are structurally positioned to corner markets; by lowering household operating costs, they naturally offset the mobility shock and expand long-run demand.

Building on the paradox framework, we introduce the **Production-Inference Tax (PIT)**—a narrowly scoped, non-distortionary 1% tax on production-level AI inference—administered

through a real-time, distributed-ledger compliance layer. Using projections from the AI market and the mobility model, we quantify how a 1% PIT offsets the fiscal damage created by the same 1% downward-mobility shock. This architecture enables verifiable inference tracking without impeding innovation while providing a stable fiscal anchor for reinforcing R2R, ESIL, and WRR at scale.

The findings illustrate that a 1% PIT, enforced via automated ledger mechanisms, generates meaningful revenue with minimal friction and strengthens economic resilience in an AI-intensive economy. Resolving the AI Productivity Paradox ultimately requires ensuring that productivity gains—whether via private-sector price transmission or public-sector fiscal design—flow back to the households whose stability sustains aggregate demand.

Finally, the paper argues that AI adoption is not optional; it is a core enabler of individual empowerment. The risk lies not in AI replacing people, but in economies failing to design systems that allow productivity gains to fortify the very households that sustain demand.

The author received no external research support for this project. The author declares no financial relationships, affiliations, or other potential conflicts of interest that could influence the research presented in this submission.