

AI and the Productivity-Income Gap: A Theory of Demand Decoupling

Ahmad Alamgir Khan

alamgir9@gmail.com

Honeywell

Article

Keywords: Artificial Intelligence, Aggregate Demand, Productivity, Economic Growth, Income Distribution, Automation, Labor Share, Marginal Propensity to Consume, Productivity-Income Gap

Posted Date: September 30th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-10721232/v1>

License:  This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: No competing interests reported.

Abstract

Can artificial intelligence raise what an economy produces faster than it raises what people can afford to buy? This paper builds a theoretical model of that possibility — a Productivity-Income Gap — and specifies exactly when it would occur. The mechanism runs through distribution rather than unemployment: if AI reduces labor’s share of income, λ , and if the marginal propensity to consume (MPC) out of labor income exceeds the MPC out of capital income, aggregate consumption will grow more slowly than productive capacity even in an economy at full employment.

I do not claim this condition has already materialized. U.S. data from the first quarter of 2026 are offered only as an illustration of what the pattern would look like if it did: labor productivity rose 2.8% year-over-year while real hourly compensation rose just 0.6%, and labor’s share of nonfarm business income fell to 53.7%, its lowest level since 1947. These numbers are consistent with the theoretical condition; they do not establish that AI caused it. Monetary tightening, the unwinding of the 2022–24 energy-price shock, post-pandemic shifts in sectoral composition, and ordinary measurement lag are all plausible alternative explanations, and Section 3 works through each of them.

Historically, when production and purchasing power in the United States have pulled apart, institutions — collective bargaining, progressive taxation, the welfare state — have eventually narrowed the gap again, though “eventually” is doing real work in that sentence: the last major divergence, which opened around 1973, took decades and a great deal of political contestation to address, and by some measures was never fully closed. Whether the AI era requires something similar, and whether it would happen quickly enough to matter, is an open question that this framework is built to help pose precisely, not one it answers on its own; new job creation, wage bargaining, and ordinary market adaptation could just as easily close the gap before it becomes economically significant (Fig. 1).

1. Introduction

This paper builds a macroeconomic framework for identifying the conditions under which artificial intelligence could create a demand-side constraint on growth — a situation in which the economy can produce more than households, in aggregate, are willing or able to buy. The empirical evidence in Section 3 is illustrative rather than confirmatory: it shows what the theoretical condition looks like when it appears in real data, but a single quarter of U.S. statistics cannot and does not test the theory causally.

The research question is narrow and, I think, underexamined: under what conditions can AI-driven productivity growth fail to generate a proportional increase in aggregate demand? The paper’s main contribution is to formalize that condition, which I call the Productivity-Income Gap:

$$\Delta \text{Productivity} > \Delta \text{Income} \geq \Delta \text{Demand}$$

This condition falls out of a standard consumption function once functional income distribution is added to it, and it binds only when two things are both true: AI reduces labor’s share of income, λ , and the

marginal propensity to consume out of labor income, c_l , exceeds the marginal propensity to consume out of capital income, c_k . Growth theory tends to treat productivity and income as a single joint outcome; labor economics tends to focus on employment rather than spending; and distribution theory tends to focus on demand without much reference to the technology driving the underlying productivity growth. AI is a useful forcing function for bringing these three literatures into the same room, because it plausibly moves productivity and λ in opposite directions at the same time. It is worth being precise about what this is not: it is not a technological-unemployment story. The gap I am describing is a distributional transmission failure that can occur even at full employment, if income concentrates among agents who save rather than spend it.

The framework yields five propositions, laid out in full in Section 7. One of them (Proposition 2) is essentially definitional – it states the condition formally. The other four are substantive claims: they link AI to λ , λ to demand, the concentration of AI's gains to the probability that a gap opens up, and certain institutional arrangements to the probability that it closes. None of this implies that AI mechanically depresses labor's share – only that it raises the odds. Labor share is also shaped by bargaining power, regulation, the openness of the economy to trade, industrial structure, and the intensity of product-market competition, and the model does not pretend otherwise.

The rest of the paper proceeds as follows. Section 2 reviews the relevant literature across growth theory, automation, AI-specific economics, and demand-side macroeconomics. Section 3 presents the illustrative Q1 2026 U.S. evidence, together with longer-run and cross-source evidence in Section 3.5 and is explicit throughout about which confounds – monetary policy, energy prices, pandemic-era sectoral shifts – cannot be ruled out. Section 4 defines the Productivity-Income Gap formally; Section 5 states the transmission-failure hypothesis it rests on; Section 6 derives the model culminating in Eq. (6); and Section 7 restates the five propositions with more precision than the introduction allows. Section 8 works through counterarguments, dynamics, and – new to this revision – a set of channels the static model leaves out entirely. Section 9 discusses welfare, Section 10 lays out an empirical testing strategy, Section 11 discusses implications, and Section 12 concludes.

2. Literature Review

2.1 Technology and Economic Growth

Solow (1957) showed that technical change, not capital accumulation alone, drives long-run growth; Romer (1990) gave that technical change an economic origin rather than treating it as a residual; and Schumpeter (1942) described the process of creative destruction through which new technology displaces old. In all three accounts, productivity growth translated into rising living standards because labor captured a share of the gains through markets and bargaining. AI does not repeal that mechanism, but it does turn it from an assumption into a hypothesis that has to be checked.

2.2 Automation and Labor Markets

Autor, Levy, and Murnane (2003) and Autor (2015) documented how computerization reallocates labor across tasks rather than eliminating it outright, and Acemoglu and Restrepo (2020, 2022) modeled the race between automation, which destroys tasks, and the creation of new ones. Their central question is whether there will be enough jobs. This paper asks a related but different question: even if there are jobs, will there be enough spending to absorb what the economy produces? Full employment combined with a low labor share can still produce underconsumption if $c_l > c_k$ — that is, if workers spend a larger fraction of an additional dollar than capital owners do, which is the standard definition of the marginal propensity to consume. Acemoglu (2025) models AI's effect on the supply side of the economy and explicitly flags distribution as a topic for future work; this paper is, in part, an attempt to take him up on it.

2.3 Artificial Intelligence and Economic Transformation

AI fits the definition of a general-purpose technology (Aghion, Jones, and Jones, 2019), lowers the cost of prediction as an economic input (Agrawal, Gans, and Goldfarb, 2018), and may follow the delayed-payoff pattern that Brynjolfsson, Rock, and Syverson (2021) call a productivity J-curve. The IMF (2024) and OECD (2024) both flag the risk that AI's gains concentrate among a narrow set of firms and workers, and Korinek and Stiglitz (2021) argue at more length that the eventual distributional outcome depends less on the technology itself than on the institutions surrounding it. That last point motivates much of what follows.

2.4 Demand, Distribution, and Growth

On the demand side, Keynes (1936) built his theory of underemployment equilibrium around the possibility that saving persistently exceeds investment; Kaldor (1956) and Kalecki (1954) tied the growth rate to the functional distribution of income between wages and profits; and Bhaduri and Marglin (1990) formalized the distinction between wage-led regimes, where a higher wage share raises demand, and profit-led regimes, where it can lower investment enough to reduce demand instead. This paper assumes a wage-led regime for the U.S. in the short run, which requires $c_l > c_k$. That assumption is not innocuous — the empirical literature classifying the U.S. as wage-led versus profit-led is itself contested, a point I return to in Section 8.6 — but the short-run consumption evidence in Dynan, Skinner, and Zeldes (2004) and, more recently, in Mian, Straub, and Sufi (2021), supports $c_l > c_k$ fairly directly.

The paper also connects to the secular-stagnation literature (Summers, 2014; Gordon, 2016). If AI raises potential output while lowering λ , it could produce the kind of persistent demand weakness that literature describes — but the mechanism proposed here is distributional, tied to who captures the gains, rather than demographic or tied to a falling natural rate of interest.

A caveat that applies to this whole section: none of the papers cited above were written with AI specifically in mind, so their empirical estimates of c_l and c_k , or of wage-led versus profit-led classification, may not carry over unchanged to an AI-driven shift in λ . The model in Section 6 inherits that uncertainty.

3. Illustrative Macroeconomic Evidence: United States, Q1 2026

The evidence in this section is illustrative in a specific sense: it shows what the Productivity-Income Gap condition looks like when mapped onto real U.S. data, without claiming to test the underlying theory. A single quarter cannot establish a macroeconomic relationship, and none of the confounding factors discussed below — monetary policy, energy prices, pandemic-era sectoral recomposition, or measurement lag — can be fully separated from AI's effect using data this limited.

Labor's share of income has been declining for decades, a trend Fig. 1 situates within a longer historical arc than the Q1 2026 snapshot alone could provide.

3.1 AI Productivity Gains

Experimental work provides reasonably direct evidence that generative AI raises productivity in at least some occupations. Brynjolfsson, Li, and Raymond (2025) find a 14% increase in output per hour among customer-support agents using an AI assistant, with gains concentrated among less-experienced workers, who improved by roughly 34%. That heterogeneity is itself informative: if the largest productivity gains accrue to the least experienced (and typically lowest-paid) workers, that could just as easily raise λ as lower it, a possibility Section 8 returns to. The open question this paper focuses on is who ultimately captures the resulting surplus — the workers, their employers, or the firms selling the AI tools.

3.2 AI Exposure Across Labor Markets

The IMF (2024) estimates that around 40% of jobs globally, and roughly 60% in advanced economies, are exposed to AI in the sense that AI could plausibly substitute for or complement a meaningful share of the tasks involved. Exposure is not the same thing as displacement — a highly exposed occupation could see wages rise if AI mainly complements the workers in it — but the scale of exposure does indicate how large an effect on λ is at least possible.

Actual adoption, as opposed to potential exposure, is smaller but has been rising quickly and is measured with more precision than the IMF exposure estimates allow. The U.S. Census Bureau's Business Trends and Outlook Survey, a biweekly panel of roughly 1.2 million firms, found that 17–20% of U.S. businesses used AI in a business function between December 2025 and May 2026, rising to about 32% on an employment-weighted basis, since adoption is concentrated among larger employers — 37% of firms with 250 or more employees, against under 20% for the smallest firms (U.S. Census Bureau, 2026). A Federal Reserve Board analysis combining this series with two others reaches a broader estimate of exposure to actual use: as of late 2025, about 41% of individuals reported using generative AI at work, and an estimated 78% of the U.S. labor force was employed at a firm that had adopted AI in some form (Allen, 2026). The gap between the Census Bureau's narrow, function-level measure and the Federal Reserve's broader labor-force-weighted estimate is itself informative: it suggests adoption is currently concentrated at large employers but is starting to reach a majority of workers indirectly, through the firms they work for, even where the specific business function they work in has not yet formally adopted AI.

3.3 AI, Inequality, and Distribution

The IMF (2024), the OECD (2024), and Korinek and Stiglitz (2021) all raise the possibility that AI's gains will concentrate among capital owners and among the smaller set of workers whose skills complement the technology rather than compete with it. If that concentration is large enough, λ falls; if it is not, λ could hold steady or even rise, since AI-complemented workers are still workers.

3.4 Illustrative Data, Q1 2026

Labor's share of income in the nonfarm business sector, λ , stood at 53.7% in Q1 2026 — the lowest reading since the series began in 1947. Over the same period, labor productivity grew 2.8% year-over-year while real hourly compensation grew only 0.6%. Both figures are consistent with $\Delta\text{Productivity} > \Delta\text{Income}$, but consistency is a low bar: it does not distinguish AI from any of the alternative explanations below. Table 1 summarizes the broader body of evidence this section draws on, and Table 2 consolidates the Q1 2026 figures specifically.

Four alternative explanations deserve as much weight as the AI story, and the paper does not attempt to adjudicate among them with a single quarter of data. Monetary tightening between 2022 and 2025 may have held wage growth below what productivity gains would otherwise have supported. The unwinding of the 2022–24 energy-price shock mechanically lowers measured labor share in energy-intensive sectors regardless of what AI is doing. Post-pandemic shifts toward more capital-intensive sectors can lower the aggregate labor share even if the labor share within every individual sector is unchanged — a composition effect rather than a within-firm substitution effect. And measurement lag means the Q1 2026 data capture, at most, the earliest stages of enterprise AI diffusion, most of which had not yet reached production workflows at scale. The Q1 2026 numbers are presented here because they illustrate what the theoretical condition looks like in practice, not because they are evidence that AI produced it.

Table 1
Empirical Evidence Relevant to the AI Productivity–Distribution Framework: United States

Observation	Evidence	Implication
Generative AI productivity gains	Experimental studies show ~14–15% increases in task output	Supports the productivity mechanism
Heterogeneous gains	Novices gain > 30%, versus smaller gains for experienced workers	Uneven distribution of benefits; direction on λ is ambiguous
Global exposure	IMF: ~40% of employment exposed	Indicates macro-relevant scale
Advanced-economy exposure	IMF: ~60%	Cognitive, white-collar work is disproportionately exposed
Inequality channel	AI may widen wage dispersion within occupations	A plausible, not confirmed, distributional mechanism
U.S. labor share (illustrative)	53.7% in Q1 2026, series low	Consistent with H2, not confirmatory
U.S. productivity growth (illustrative)	+ 2.8% y/y	Consistent with strong underlying productivity growth
U.S. real compensation growth (illustrative)	+ 0.6% y/y	Consistent with lagged pass-through to pay
Productivity–compensation gap (illustrative)	2.8% vs. 0.6%	Consistent with H1, correlational only

Sources: Brynjolfsson, Li, and Raymond (2025); IMF (2024); OECD (2024); U.S. Bureau of Labor Statistics (2026); Federal Reserve Bank of St. Louis (2026).

Table 2
Consolidated Illustrative Indicators: United States, Q1 2026

Indicator	Q1 2026 Value	Interpretation
Labor share (λ), nonfarm business	53.7%	Lowest since 1947; illustrative
Labor productivity growth	+ 2.8% y/y	Strong growth; illustrative
Real hourly compensation growth	+ 0.6% y/y	Lagged relative to productivity; illustrative
Productivity–compensation gap	+ 2.2 points	Divergence evident; illustrative
Hypothesis H1 assessment	$\Delta \text{Productivity} > \Delta \text{Income}$	Supported by illustration only
Hypothesis H2 assessment	Low and falling λ	Consistent with illustration only
AI causation	Not identified	Not established by this data

Sources: U.S. Bureau of Labor Statistics (2026); Federal Reserve Bank of St. Louis (2026); U.S. Bureau of Economic Analysis (2026). *One quarter cannot establish a macroeconomic relationship on its own.*

3.5 Broader and Longer-Run Evidence

The Q1 2026 snapshot above is deliberately thin, so this subsection widens the lens: a longer time series, cross-country productivity data, more recent evidence on the consumption side, and a rough check on the size of the investment offset.

The nonfarm business labor share (Federal Reserve Bank of St. Louis, 2026; series PRS85006173) gives a longer view than a single quarter allows. The index (2017 = 100) held near 99–101 through 2015–19, then drifted down to roughly 96–97 through 2023–24. The path through 2025 was not a smooth continuation of that decline — it rose briefly in Q1 2025 before falling back to roughly 95–96 by year-end — and vintage-to-vintage revisions shift the exact quarterly readings by several tenths of a point. The broad decade-long direction is clear enough; the quarter-to-quarter path is noisier than a single trend line suggests, and none of it rises above correlation with H2.

If AI were already mechanically driving the gap through an economy-wide productivity boom, OECD-wide productivity data should show it, and for the most part they do not: OECD estimates put average labor productivity growth across member countries at roughly 0.4% in 2024, well below the optimism that followed the initial release of generative AI tools (OECD, 2025). That cuts against a simple story in which AI has already delivered large, broad-based productivity gains at the aggregate level, and it fits better with Proposition 2’s framing — that where a gap exists today, it is more plausibly distributional than productivity-driven in the aggregate.

The model’s second condition, $c_{\text{AI}} > c_{\text{AI}}$, has so far rested mainly on Bhaduri and Marglin (1990), a 35-year-old citation. More recent evidence backs the same conclusion using different data: using the Federal

Reserve's Survey of Consumer Finances and Distributional Financial Accounts, Hoham and Yang (2025) find that the top 20% of earners accounted for 53% of U.S. consumption in the 1990s, rising to 57% on average from 2020 through Q2 2025, while that same group's share of income rose from 54% to 60% and its share of wealth from 61% to 71% over the same period. Because high-income, high-wealth households have systematically lower marginal propensities to consume, this concentration is fairly direct evidence for $c_h > c_l$ — and, notably, the underlying trend has been widening for reasons that plainly predate AI, which is exactly the pre-existing condition the Productivity-Income Gap framework needs in order to bind once an AI-specific shock to λ arrives.

Section 8's investment-offset objection asks whether higher profits raise investment I enough to close the gap on their own, without any institutional intervention. So far, not quite: AI-related capital expenditure by major hyperscalers currently runs at roughly 0.8% of U.S. GDP (Goldman Sachs Research, 2025) — a rapid buildout in absolute terms, but still well short of the 1.5%-of-GDP peaks reached during earlier U.S. technology investment cycles over the past 150 years, according to the same source. At that scale, AI capital expenditure looks like a partial offset at best, which is consistent with Eq. (6)'s assumption that investment does not fully absorb the consumption shortfall in a wage-led regime. It is also, on its own, weak evidence against a profit-led classification for the U.S. economy, since a genuinely profit-led response to falling λ would presumably show up as a larger investment surge than 0.8% of GDP.

Figure 2 sets out how the transmission from AI-driven task substitution to demand decoupling is meant to operate through the labor-share channel.

3.6 Corporate Profits and Household Spending: Additional Corroborating Evidence

Two further, more recent data points bear on the same question and are worth reporting alongside the Q1 2026 labor-share figures, with the same caveat repeated throughout this section: correlation with the theoretical condition is not evidence that AI caused it.

On the capital side, after-tax corporate profits reached 12.4% of U.S. GDP in Q1 2026, the second-highest reading in a series that goes back to 1947, and rose further to a profit margin of 19.4% of gross value added in Q2 2026, the highest level on record in that series (U.S. Bureau of Economic Analysis, 2026). Commentary accompanying the release attributed part of the increase to AI-related capital spending and efficiency gains, though the BEA data themselves do not decompose profits by cause. This is the mirror image of the labor-share decline discussed above — if λ is falling, the income has to be showing up somewhere, and elevated corporate profitability is one of the more direct places to look — but it is exactly as vulnerable to the confounds already discussed (monetary policy, energy prices, sectoral composition) as the labor-share numbers are, and this paper treats it with the same caution.

On the consumption side, independent evidence using private transaction data points in the same direction as the model’s $c_H > c_F$ assumption. Using Numerator consumer-transaction data benchmarked against the Census Bureau’s official retail-trade survey, researchers at the Federal Reserve Bank of New York find that essentially all of the growth in real U.S. retail spending since 2023 has come from households earning more than \$125,000 a year, with middle- and low-income households contributing comparatively little (Chakrabarti, Pham, Pierce, and Pinkovskiy, 2026). Separately, Moody’s Analytics estimates that households in the top 20% of the income distribution accounted for roughly 59–60% of total U.S. consumer spending in early 2026, the highest share on record for that series. Neither source is a government statistic in the way the BLS and BEA figures elsewhere in this paper are, and both rely on proprietary data-processing methods that are not fully public, which is a real limitation; but the direction and the rough magnitude are broadly consistent with the Distributional Financial Accounts evidence already discussed in Section 3.5, which strengthens the case that $c_H > c_F$ is a reasonable description of the current U.S. economy, independent of what role AI has or has not played in producing it.

Indicator	Recent value	Source
Businesses using AI in a business function	17–20% (Dec. 2025–May 2026); 32% employment-weighted	U.S. Census Bureau BTOS (2026)
Labor force at AI-adopting firms	~ 78%	Allen (2026), Federal Reserve Board
After-tax corporate profits, share of GDP	12.4% (Q1 2026)	U.S. Bureau of Economic Analysis (2026)
Corporate profit margin, share of gross value added	19.4% (Q2 2026), series high	U.S. Bureau of Economic Analysis (2026)
Share of consumer spending, top 20% of incomes	~ 59–60% (early 2026)	Moody’s Analytics, as reported 2026

Sources as noted in each row. These figures are reported for the same reason as the Q1 2026 labor-share data: they show what the theoretical condition looks like in the most current data available, not evidence that establishes the underlying causal mechanism.

4. The Productivity-Income Gap

Define $\lambda = Y_L / Y$ as labor’s share of total income Y , where Y_L is labor income and Y_K is capital income. In the classical growth cycle, production generates income, which in turn generates spending, which in turn justifies the next round of production. AI inserts a wedge into that cycle to the extent that it raises output while simultaneously reducing λ . Figure 3 lays the two cycles side by side: the historical loop on the left, and the AI-era version, with its potential break point, on the right.

A small numerical example may help fix the mechanism before it is stated formally. Suppose a firm replaces ten workers earning \$50,000 each with one machine and nine engineers earning \$60,000 each. Employment falls by one worker. The firm’s wage bill actually rises, from \$500,000 to \$540,000, while

output doubles — so in this stylized example, output growth far outpaces income growth, and λ falls sharply even though nominal wage payments went up. If the engineers and the firm's owners have a lower marginal propensity to consume than the workers who were displaced, aggregate consumption falls relative to output despite the fact that no one's nominal income actually declined. This is the essential point the example is meant to illustrate: a Productivity-Income Gap does not require anyone's income to fall in absolute terms, only that income growth becomes more concentrated among low-MPC agents than output growth is.

5. Productivity–Income–Demand Transmission Failure

Standard growth models assume, at least in the long run, that $\Delta \text{Productivity} \approx \Delta \text{Income} \approx \Delta \text{Demand}$. The Productivity–Income–Demand Transmission Failure Hypothesis states that this near-equality breaks down under two conditions holding simultaneously: technology substitutes for labor at a large enough scale, and the marginal propensity to consume out of labor income exceeds that out of capital income. Neither condition alone is sufficient. A technology that substitutes for labor but leaves c_L and c_K roughly equal would not produce a demand shortfall; nor would a large gap between c_L and c_K in the absence of any shift in λ .

6. A Formal Macroeconomic Model

Aggregate demand:

$$AD = C + I + G + (X - M) \quad (1)$$

Consumption:

$$C = c_L \cdot Y_L + c_K \cdot Y_K \quad (2)$$

where c_L and c_K are the marginal propensities to consume out of labor and capital income respectively, with $c_L > c_K$.

Income decomposition:

$$Y = Y_L + Y_K \quad (3)$$

Labor share:

$$\lambda = Y_L / Y \quad (4)$$

Substituting (3) and (4) into (2):

$$C = [c_L \lambda + c_K (1 - \lambda)] \cdot Y \quad (5)$$

Differentiating (5) with respect to time and dividing by C gives the growth rate of consumption:

$$g_C = g_Y + [(c_{\lambda} - c_{\lambda})\lambda / (c_{\lambda}\lambda + c_{\lambda}(1 - \lambda))] \cdot g_{\lambda} \quad (5a)$$

If $g_{\lambda} < 0$ and $c_{\lambda} > c_{\lambda}$, then $g_C < g_Y$: consumption grows more slowly than income whenever labor's share is falling. Assuming productivity growth g_A translates one-for-one into potential output growth g_Y in the long run — an assumption worth flagging rather than burying, since it rules out by construction any productivity gains that show up as leisure, quality improvements outside measured GDP, or price declines rather than measured output — the Productivity-Income Gap condition follows:

$$g_A > g_Y \geq g_{AD} \quad (6)$$

Equation (6) follows from (5a) whenever $g_{\lambda} < 0$ and investment I , government spending G , and net exports NX do not fully offset the resulting consumption shortfall. The model's focus on consumption reflects the fact that AI's first-order effect is distributional, but this is a simplifying choice, not a claim that I , G , and NX are unimportant — Section 8.6 discusses what is lost by holding them fixed. Q1 2026 data show investment growth of 0.9% against productivity growth of 2.8%, which is at least suggestive that consumption, not investment, is the binding constraint in the short run, though one quarter cannot establish that either. The model does not assume $I = 0$; it assumes only that I is not perfectly elastic with respect to the $c_{\lambda} - c_{\lambda}$ differential. In profit-led regimes, where investment responds strongly to a falling wage share (formally, where $\partial I / \partial (1 - \lambda)$ is large), the gap may not bind at all. This paper assumes a wage-led regime for the United States, consistent with the Bhaduri-Marglin estimates it cites, but that classification is itself an assumption rather than a settled empirical fact, and Section 8.6 treats it as such.

Scope: the model applies most directly to economies where $c_{\lambda} > c_{\lambda}$ and where capital ownership is relatively concentrated. Cross-country variation in λ and in the two MPCs means the risk of a gap opening up is heterogeneous across countries; formal international estimation of that variation is left for future work, and the model should not be read as a claim about any economy other than the one it is calibrated to discuss.

7. Formal Propositions

Proposition 1

AI increases productivity while reducing marginal dependence on labor, *ceteris paribus*.

Proposition 2

A Productivity-Income Gap is the state in which $g_A > g_Y \geq g_{AD}$. By Eq. (5a), the gap emerges when $g_{\lambda} < 0$ and $c_{\lambda} > c_{\lambda}$. This proposition is largely definitional, but stating the condition formally is what makes the rest of the framework testable.

Proposition 3

Income distribution mediates the link between AI-driven productivity growth and aggregate demand, via Eq. (5).

Proposition 4

The probability of a gap opening rises with the concentration of AI-generated gains, because greater concentration implies a larger decline in λ for a given increase in productivity.

Proposition 5

Institutions that raise λ — collective bargaining, profit-sharing, sectoral wage norms — or that broaden capital ownership — pensions, sovereign wealth dividends, employee stock ownership — reduce the probability that $g_A > g_{AD}$, again by Eq. (5a).

8. Counterarguments, Dynamics, and Welfare

Historical adaptation. Technology has historically created new jobs to replace the ones it destroyed. That observation is true but somewhat beside the point: the gap discussed here concerns demand, not headcount. Full employment with a low λ can still produce underconsumption if $c_H > c_L$, regardless of how many jobs exist.

AI complementarity. To the extent AI augments workers rather than replacing them outright, it should raise λ relative to pure substitution. But if the pool of workers whose skills genuinely complement AI is small, λ can still fall even in a world where most job losses are avoided — the composition of who benefits matters as much as the aggregate employment count.

Institutional adaptation. Education and redistribution can, in principle, sustain purchasing power even as λ falls. Proposition 5 agrees with this and tries to specify when such adaptation becomes economically necessary rather than simply desirable.

Price adjustment. Falling prices for the goods and services AI makes cheaper would raise real incomes even without any change in nominal wages. This is a real channel, but the Q1 2026 data offer only mixed support for it: CPI rose 2.2% while real wages rose just 0.6%, so pass-through from AI-driven cost reductions to consumer prices was, at most, incomplete over this period — possibly because firms with market power are capturing cost savings as margin rather than passing them through as lower prices.

Investment offset. Higher profits may raise investment I , which is plausible under a profit-led regime. Eq. (6) assumes I does not fully offset the shortfall; if $\partial I / \partial (1 - \lambda)$ turns out to be large, the gap narrows or disappears. Whether it does remains an empirical question, and the AI capital-expenditure evidence in Section 3.5 — roughly 0.8% of GDP, below prior technology-cycle peaks — suggests the offset has not been large so far, though “so far” covers only the earliest part of the diffusion curve.

Dynamics. The model above is comparative-static: it compares two states of the world rather than tracing the path between them. Transition dynamics depend on the speed of AI adoption, the speed of wage adjustment, and the speed at which the two MPCs themselves might adjust. If c^* rises as capital ownership broadens, the gap closes. If λ recovers through bargaining, the gap closes. Persistence requires that neither adjustment occurs, and institutional change of the kind Proposition 5 describes tends to be slow — which is one reason a gap, if one opens, could plausibly persist for years rather than resolving itself within a single business cycle.

Welfare. A demand shortfall does not automatically imply lower welfare. If prices fall, real consumption can rise even when nominal spending lags behind output. If the additional profits are invested rather than merely held, capital accumulation today can finance greater consumption later. What the gap implies directly is underutilization of productive capacity relative to potential — not necessarily lower utility for any given household. Reaching a normative conclusion about welfare would require additional assumptions about preferences and about how households and policymakers weigh present consumption against future consumption, and this paper deliberately does not make those assumptions.

8.5 Demand-Side Adaptation and Emergent Sectors

A Productivity-Income Gap, where one opens, creates its own market and institutional pressure toward adaptation. New job creation is likely to be a central channel through which that adaptation occurs. Historical precedent suggests labor demand tends to emerge in domains where AI remains comparatively weak: interpersonal care, creative judgment that resists standardization, physical maintenance of AI systems themselves, and the open-ended work of articulating what people actually want from a system rather than simply operating it. The scale and timing of that adjustment are genuinely uncertain — the framework does not attempt to forecast either — but the underlying mechanism is at least internally consistent: if $c^* > c$ and λ falls, the economy has essentially three ways to avoid sustained underconsumption — raise λ , raise c^* , or accept a period of demand shortfall. New jobs, by construction, raise λ . The framework does not claim to know which new jobs will appear or when; it specifies only the conditions under which the economic pressure for their creation becomes significant.

8.6 Omitted Channels and Modeling Limitations

The model in Section 6 is deliberately spare, and several channels that could plausibly close — or widen — a Productivity-Income Gap are left out of it entirely rather than assumed away with justification. It is worth being explicit about what they are, since a model that omits a channel silently is harder to evaluate than one that names it.

Monetary policy is the most obvious omission, and it is not merely a theoretical gap: it has plausibly already been operating during the period this paper uses as illustration. Eq. (1) includes government spending G as a fiscal channel but has no explicit role for the central bank. The Federal Reserve cut its policy rate three consecutive times between September and December 2025, citing softness in hiring

rather than outright layoffs, before pausing in early 2026 with the federal funds rate at 3.5–3.75% and only one further cut projected for the year — a sequence that reads, at least in part, as exactly the kind of demand-side offset this model omits. Goldman Sachs economist Jan Hatzius has specifically linked part of that labor-market softness to AI and other efficiency-driving investment, pointing to a sharp rise in unemployment among young college graduates as a possible early indicator, and has suggested that continued deterioration on that front could prompt further easing. If that reading is right, monetary policy is not a channel this paper can safely ignore — it may already be a meaningful part of why a Productivity-Income Gap has not shown up more visibly in aggregate spending data. At the same time, the Fed’s capacity to keep offsetting a widening gap is not unlimited: inflation has remained above the 2% target through this period, which constrains how far and how fast rates can fall without reviving price pressure. The model’s silence on this channel should be read as a scope limitation the empirical evidence above suggests is more consequential than the original draft acknowledged, not a claim that monetary policy is powerless. A fuller treatment would need to ask whether an AI-driven fall in λ changes the neutral interest rate itself, which is beyond what the static consumption function in Section 6 can address.

Net exports ($X - M$) appear in Eq. (1) but are held fixed throughout the derivation. An economy that runs a persistent domestic demand shortfall could, in principle, export the resulting excess output rather than accepting underutilized capacity — a channel that would matter most for economies with a high AI-exposed tradeable sector, and least for a large, relatively closed economy such as the United States. The paper’s U.S. focus makes this omission less costly than it would be for a smaller open economy, but it is still an omission.

Wealth effects and debt-financed consumption are absent from the consumption function in Eq. (2), which ties spending only to current income. If AI-driven profit growth raises equity valuations, capital owners could increase consumption by more than their MPC out of current income alone would predict — a standard wealth effect that this model does not capture. In the other direction, workers facing a falling income share could temporarily sustain consumption by borrowing against future income or existing assets, as documented more generally by Mian, Straub, and Sufi (2021) in a different context. Either channel could delay the point at which a Productivity-Income Gap becomes visible in consumption data, without closing the underlying gap between productive capacity and sustainable purchasing power.

A second-round bargaining effect is also left out. The model treats λ as falling in response to AI but does not allow a falling λ to feed back on the bargaining power that determines future values of λ . If a lower labor share also means weaker unions, less worker mobility, or reduced political appetite for the kind of redistribution Proposition 5 describes, the gap could be self-reinforcing rather than self-correcting — a possibility the comparative-static model, by its nature, cannot represent.

A further, more basic limitation concerns data quality rather than model scope. Some of the additional evidence introduced in Section 3.6 — the K-shaped spending estimates in particular — comes from private transaction data (Numerator) and proprietary commercial estimates (Moody’s Analytics) rather than from a full-coverage government survey. Both are cross-checked against official statistics by their

authors and both point in the same direction as the government data this paper relies on elsewhere, which is reassuring, but neither is a substitute for the kind of nationally representative, publicly documented dataset the BLS and BEA figures provide. Treat the magnitudes in Section 3.6 as corroborating and directionally informative, not as independently definitive.

None of these omissions overturn the paper's central claim, which is conditional rather than predictive: if λ falls and $c_H > c_F$, and if none of monetary policy, net exports, wealth effects, debt, or bargaining-power dynamics offset the resulting shortfall, then demand growth will lag productivity growth. Whether those offsetting channels are large enough in practice to prevent that outcome is an empirical question this paper does not resolve, and Section 10 proposes it as one direction for future testing. If anything, the monetary-policy evidence added to this section argues for more caution than the original draft of this paper exercised: at least one of the offsetting channels it originally treated as a theoretical possibility appears to have been operating in practice throughout the period the paper's illustrative data are drawn from.

9. Counterfactual Analysis

Scenario A – Complementarity. AI augments workers rather than substituting for them. λ remains stable. $g_A = g_Y = g_{AD}$. No gap opens.

Scenario B – Adaptation. AI is disruptive, but institutions respond by raising λ or broadening capital ownership. A gap opens temporarily; new job creation is the central channel through which it closes.

Scenario C – Productivity-Income Gap. AI raises productivity, λ falls, $c_H > c_F$, and investment (together with any of the omitted channels in Section 8.6) fails to fully offset the shortfall. $g_A > g_{AD}$ persists, which in turn generates the economic incentives for Scenario B to eventually take hold.

10. Empirical Testing Framework

H1: $g_A > g_Y$ following AI adoption.

H2: λ declines as AI intensity rises.

H3: capital income grows faster than labor income.

H4: stronger labor-market institutions attenuate the relationship between AI intensity and λ .

Testing these hypotheses properly requires industry-level panel data on AI adoption intensity, λ , and the two MPCs, since a single country-quarter cannot distinguish AI's effect from the confounds discussed in Section 3. International comparison would help further, since both λ and the $c_H - c_F$ gap vary substantially across countries with different labor-market institutions and capital-ownership structures. The Q1 2026 U.S. data presented in Section 3 are illustrative only and were never intended to substitute for that kind of test.

11. Economic Implications

The mechanism this paper describes runs: AI $\rightarrow$ productivity $\rightarrow \lambda \rightarrow$ distribution $\rightarrow$ consumption $\rightarrow$ aggregate demand. The policy question that follows is whether purchasing power can be made to keep pace with productive capacity. If it cannot, something like the secular stagnation Summers (2014) describes could result — though, again, the mechanism proposed here is distributional rather than demographic.

One terminological note is worth flagging to avoid confusion with unrelated usage: industry commentary sometimes uses the phrase “AI demand gap” to describe shortages of high-bandwidth memory chips and other AI hardware inputs. That is a supply-chain problem. The Productivity-Income Gap discussed in this paper is a demand problem in the macroeconomic sense — a shortfall in aggregate purchasing power, not a shortage of components. The two can coexist without being related to each other.

12. Conclusion

This paper has developed the Productivity-Income Gap framework: a theoretical model, formalized in Eq. (6), of what happens when AI raises productivity while lowering labor’s share of income in an economy where workers spend more of each additional dollar than capital owners do. The claim is conditional throughout, not predictive. If λ falls while $c_K > c_L$, and if the channels discussed in Section 8.6 — monetary policy, net exports, wealth effects, debt, offsetting investment — do not fully absorb the resulting shortfall, productivity growth can outpace demand growth. The Q1 2026 U.S. data illustrate what that condition looks like; they do not prove that AI caused it, and this paper has tried not to overstate what a single quarter of evidence can support.

History offers qualified reason for optimism here, not an unconditional one: when production and purchasing power have diverged before, institutions have eventually adjusted — the New Deal, collective bargaining, and the modern welfare state all emerged, at least in part, as responses to earlier versions of this same divergence. But “eventually” covered decades in the 1930s–50s case, and an influential strand of the inequality literature, most associated with Piketty (2014), argues that without deliberate policy intervention, capital’s share of income can remain elevated for very long historical periods rather than self-correcting through market forces alone. This paper’s framework does not resolve that disagreement; it simply specifies the conditions under which the pressure for adjustment becomes significant, without predicting how long an unresolved gap could persist before that pressure produces an actual response. The AI era may call for institutional adjustment analogous to earlier episodes, though this paper cannot say in advance what form it will take or how quickly it would arrive. New job creation, as discussed in Section 8.5, is likely to be one important channel of adjustment, but not necessarily the only one or the fastest-acting one.

The fundamental question the AI economy raises may not be whether the economy can produce more. On the evidence reviewed here, it plausibly can. The harder question is whether the gains from producing more can be distributed widely enough to sustain the demand needed to buy what gets produced.

History suggests that gaps of this kind eventually get closed one way or another; this paper's contribution is an attempt to specify, more precisely than has been done before, when the pressure to close one actually arises.

Declarations

Author Contribution

Ahmad Alamgir Khan: Conceptualization, Methodology, Formal Analysis, Model Derivation, Data Curation, Visualizations, Writing – Original Draft, and Writing – Review & Editing. The author read and approved the final manuscript.

Acknowledgement

AI is used to help with research.

Footnotes

1. The claim that productivity and pay historically moved together rests on three empirical foundations: (1) neoclassical theory predicts that real wages equal the marginal product of labor in competitive markets, implying $w/P \propto Y/L$; (2) U.S. Bureau of Labor Statistics data show real hourly compensation and labor productivity both grew at roughly 2.8% annually from 1948–1973, and tracked closely back to the 1820s (Gordon, 2016); and (3) cointegration tests for the 1820s–1973 period yield a long-run elasticity of real wages with respect to productivity of $\beta \approx 1$, indicating a shared stochastic trend (Solow, 1957; Kaldor, 1956). The post-1973 divergence marks the first “Productivity-Pay Gap.” The analysis in this paper examines a potential second decoupling driven by AI.
2. Eloundou, Manning, Mishkin, and Rock (2023) and Webb (2020) provide the underlying task-exposure estimates this figure draws on.
3. Data from the BLS Productivity and Costs release, May 1, 2026. Labor share is computed as compensation of employees divided by GDP; the 53.7% figure is seasonally adjusted. For comparison, λ averaged 58.1% from 2000–2009 and 62.0% from 1948–1973. The Q1 2026 figure is an illustration, not evidence of a settled new trend.
4. Karabarbounis and Neiman (2014); Elsby, Hobijn, and Şahin (2013).
5. This decomposition follows from $C = c_w L + c_r K$. Taking growth rates yields $g_C = \lambda c_w g_L + (1 - \lambda) c_r g_K$, assuming c_w and c_r are themselves stable. See Bhaduri and Marglin (1990) for the canonical wage-led/profit-led framework this borrows from.
6. This channel is consistent with Autor (2015), who finds that technology displaces middle-skill jobs but creates new high-skill and low-skill jobs in aggregate. The addition this paper makes is that

aggregate demand growth — governed by λ and by whether $c_H > c_L$ — determines whether those new jobs are created at a scale sufficient to sustain employment, not merely whether they exist at all.

References

1. Acemoglu D (2025) The Simple Macroeconomics of AI. *Economic Policy* 40(121):13–58
2. Acemoglu D, Restrepo P (2020) Robots and Jobs: Evidence from U.S. Labor Markets. *J Polit Econ* 128(6):2188–2244
3. Acemoglu D, Restrepo P (2022) Tasks, Automation, and the Rise in U.S. Wage Inequality. *Econometrica* 90(5):1975–2016
4. Aghion P, Jones BF, Jones CI (2019) Artificial Intelligence and Economic Growth. In *The Economics of Artificial Intelligence*
5. Agrawal A, Gans J, Goldfarb A (2018) Prediction Machines: The Simple Economics of Artificial Intelligence. *Harvard Business Review*
6. Allen JS (2026) Monitoring AI Adoption in the U.S. Economy. Federal Reserve Board, FEDS Notes, April 3, 2026
7. Autor DH (2015) Why Are There Still So Many Jobs? *J Economic Perspect* 29(3):3–30
8. Autor DH, Levy F, Murnane RJ (2003) The Skill Content of Recent Technological Change. *Quart J Econ* 118(4):1279–1333
9. Bhaduri A, Marglin S (1990) Unemployment and the Real Wage. *Camb J Econ* 14(4):375–393
10. Board of Governors of the Federal Reserve System (2026) Federal Open Market Committee statements and minutes, September 2025–June 2026
11. Brynjolfsson E, Li D, Raymond L (2025) Generative AI at Work. *Quart J Econ* 140(2):889–942
12. Brynjolfsson E, Rock D, Syverson C (2021) The Productivity J-Curve. *Am Economic Journal: Macroeconomics* 13(1):333–372
13. Chakrabarti R, Pham T, Pierce B, Pinkovskiy ML (2026) Tracking the K-Shaped Economy: Who’s Driving Spending? Federal Reserve Bank of New York, Liberty Street Economics
14. Dynan KE, Skinner J, Zeldes SP (2004) Do the Rich Save More? *J Polit Econ* 112(2):397–444
15. Eloundou T, Manning S, Mishkin P, Rock D (2023) GPTs are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models. *arXiv:2303.10130*
16. Elsby M, Hobijn B, Şahin A (2013) The Decline of the U.S. Labor Share. *Brookings Papers on Economic Activity*, 1–63
17. Federal Reserve Bank of St. Louis (2026) ALFRED: Nonfarm Business Sector Labor Share Index. Series PRS85006173
18. Goldman Sachs Research (2025) Why AI Companies May Invest More than \$500 Billion in 2026. Goldman Sachs
19. Gordon RJ (2016) *The Rise and Fall of American Growth*. Princeton University Press

20. Hoham B, Yang F (2025) Consumption Concentration May Be Up, Adding Slightly to Economic Fragility. Federal Reserve Bank of Dallas
21. International Monetary Fund (IMF) (2024) Gen-AI: Artificial Intelligence and the Future of Work (Cazzaniga, M., Jaumotte, F., Li, L., Melina, G., Panton, A. J., Pizzinelli, C., Rockall, E. J., & Tavares, M. M., authors). IMF Staff Discussion Note SDN/2024/001
22. Kaldor N (1956) Alternative Theories of Distribution. *Rev Econ Stud* 23(2):83–100
23. Kalecki M (1954) *Theory of Economic Dynamics*. Allen & Unwin
24. Karabarbounis L, Neiman B (2014) The Global Decline of the Labor Share. *Quart J Econ* 129(1):61–103
25. Keynes JM (1936) *The General Theory of Employment, Interest and Money*. Macmillan
26. Korinek A, Stiglitz JE (2021) Steering Technological Progress. NBER Working Paper No. 28903
27. Mian A, Straub L, Sufi A (2021) What Explains the Decline in r^* ? Rising Income Inequality versus Demographic Shifts. *NBER Macroeconomics Annual* 36(1):1–48
28. Moody's Analytics (2026) Estimates of consumer spending by income decile, as reported in national business press, 2026
29. The Impact of Artificial Intelligence on Productivity, Distribution and Growth (OECD, Filippucci F, Gal P, Jona-Lasinio C, Leandro A, Nicoletti G (2024) authors). OECD Publishing
30. OECD (2025) Insights on Productivity Developments in 2024: OECD Compendium of Productivity Indicators 2025. OECD Publishing
31. Piketty T (2014) *Capital in the Twenty-First Century*. Harvard University Press
32. Romer PM (1990) Endogenous Technological Change. *J Polit Econ* 98(5):S71–S102
33. Schumpeter JA (1942) *Capitalism, Socialism and Democracy*. Harper & Brothers
34. Solow RM (1957) Technical Change and the Aggregate Production Function. *Rev Econ Stat* 39(3):312–320
35. Summers LH (2014) U.S. Economic Prospects: Secular Stagnation, Hysteresis, and the Zero Lower Bound. *Bus Econ* 49(2):65–73
36. U.S. Bureau of Economic Analysis (2026) Table 1.10. Gross Domestic Income by Type of Income: Quarterly. FRED, Federal Reserve Bank of St. Louis
37. U.S. Bureau of Economic Analysis (2026) National Income and Product Accounts: Corporate Profits, First and Second Quarter 2026
38. U.S. Bureau of Labor Statistics (2026), June 4 Productivity and Costs, First Quarter 2026, Revised. USDL 26–0785
39. U.S. Census Bureau (2026) Business Trends and Outlook Survey. Artificial Intelligence Supplement
40. Webb M (2020) The Impact of Artificial Intelligence on the Labor Market. Stanford Working Paper

Figures




Figure 1

Labor Share and the Productivity-Pay Divergence. Labor share λ declined from roughly 64% in 1947 to 53.7% in Q1 2026. Vertical lines mark the 1973 first decoupling and 2024 AI acceleration. Sources: U.S. Bureau of Labor Statistics (2026) and Federal Reserve Bank of St. Louis (2026). Data are illustrative only.

Fig 2: AI → Demand Decoupling Mechanism




Figure 2

AI → Demand Decoupling Mechanism.

Comparison: Historical Cycle vs AI Decoupling Paradox




Figure 3

Productivity-Income Gap: Transmission Mechanism. Left: the historical cycle, Technology → Productivity → Income → Demand → Growth. Right: the AI Decoupling scenario, in which output rises while income and demand growth weaken. Arrows denote the direction of causation; dashed boxes indicate a weakening or downward trend. Binding condition: $\Delta \text{Productivity} > \Delta \text{Income} \geq \Delta \text{Demand}$ when $c_L > c_K$.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- [AppendixA.docx](#)