Discussion of "Human Capital as a Driver of Automation: Theory and Evidence from China's College Expansion"*

Zhang Chen[1]

Running head: Discussion of Chen et al. (2026)

JEL classification: J24, E25, O33, O41

Keywords: Automation; Human capital; Labor share; Economic growth

1 Background

While artificial intelligence is undeniably one of the most prominent topics today, automation has been transforming the production of goods and services for centuries. The mechanization of agriculture, the rise of factory production during the Industrial Revolution, and the introduction of computers and robotics in manufacturing are just a few examples of how automation has reshaped economic structures and labor markets over time. A defining feature of automation is that capital substitutes for labor in performing production tasks. However, this observation raises an important puzzle. Despite repeated waves of industrialization and technological change, the aggregate labor share remained relatively stable for much of the twentieth century and only began to decline in recent decades. This stability suggests the presence of countervailing forces that offset the labor-reducing effects of automation.

This puzzle has motivated two prominent theoretical frameworks for understanding automation and its macroeconomic implications. Acemoglu and Restrepo (2018) propose a framework based on the creation of new tasks. Although automation replaces labor in some tasks, new tasks and occupations are created at the same time, restoring labor demand and helping stabilize the labor

*First submitted on July 20, 2026; revised on August 30, 2026.

[1]Department of Economics, Hong Kong University of Science and Technology, Clear Water Bay, Kowloon, Hong Kong. zhangchen.econ@gmail.com.

share. An alternative framework, developed by Aghion et al. (2019) and Jones and Liu (2024), draws on the logic of Baumol's cost disease. As tasks become automated, their unit costs fall, which reduces their value-added share in aggregate output. This creates a stabilizing force that counteracts the downward pressure of automation on the labor share, even when the share of automated tasks approaches one.

Inspired by these theoretical insights, a growing literature has examined the drivers of automation and its consequences for economic growth, labor markets, and the income distribution. Restrepo (2024) provide a comprehensive review of these developments. The paper under discussion contributes to this literature by applying this line of theory to the context of China.

China's recent performance in automation deserves particular attention. In recent years, the country has emerged as a major player in the new wave of artificial intelligence and automation, with developments such as the rise of AI firms and advances in humanoid robotics drawing global attention. These developments are unlikely to be accidental; rather, they may reflect the long-run effects of policy decisions and strategic investments made years earlier. One of the most consequential policy initiatives in China's recent history was the 1999 expansion of higher education, which substantially increased the supply of skilled labor starting in 2003, when the first expanded cohort graduated.

The paper under discussion seeks to link college expansion to China's rapid growth and technological upgrading. More specifically, it asks whether part of China's growth experience can be explained by increased automation, itself driven in part by the expansion in the supply of skilled labor. This question is important for understanding the rapid development of automation technologies in China, and it also has broader implications for understanding how human capital accumulation can shape the trajectory of automation and growth in other settings. Overall, I view the paper as a valuable contribution to the literature on automation, growth, and structural change. The comments below highlight what I see as the paper's main contributions and place its results in the broader context of automation and labor-share dynamics.

2 Contribution

In my view, the paper's main contribution is to connect a modern task-based model of automation to a concrete and important macroeconomic episode in China. It provides both theory and evidence on the relationship between skilled labor supply and automation-led growth. On the theory side, the paper extends the framework in Jones and Liu (2024) by introducing heterogeneous labor. While unskilled labor is used only in production, skilled labor can be allocated across three activities: production, vertical innovation, and horizontal innovation (automation). This setup gives the model a flexible mechanism through which an increase in the supply of skilled labor can have distinct effects on automation, innovation, and output, making the framework well suited for empirical analysis.

On the empirical side, the paper uses China's 1999 college expansion as a plausibly exogenous shift in the supply of skilled labor and examines whether this shock increased automation, measured primarily through robot adoption. This is a natural empirical setting for the question, given that China in the early 2000s experienced both a major expansion in higher education and rapid changes in production technology. Specifically, the paper identifies the impact of the college expansion on automation through cross-industry variation in the use of skilled labor, providing micro-level evidence on the relationship between skilled labor and automation.

More broadly, I see the paper as especially valuable because it attempts to bring the theory in Jones and Liu (2024) to the data. Given the parsimony of the original Jones–Liu framework, such quantification exercises are important for assessing the empirical relevance of the underlying mechanisms and for clarifying which margins may be most promising for future research. In this sense, the quantitative results in the paper are informative about future avenues for research on automation. In the following section, I place the paper's results in a broader context and discuss their implications.

3 Putting the Paper's Results into a Broader Context

In the Jones–Liu framework, the aggregate labor share is shaped by two countervailing forces: the downward pressure from automation and the upward pressure from capital-augmenting quality improvements. Following their notation, let $\beta_t \in [0, 1]$ denote the share of automated tasks and let Z_t denote the harmonic average of capital-augmenting productivities. The aggregate state variable is $s_t \equiv \beta_t Z_t^{-1}$, and the aggregate labor share is a decreasing function of s_t. Greater automation (a higher β_t) lowers the aggregate labor share, whereas quality improvement (a higher Z_t) raises it. Balanced growth occurs when Z_t and β_t grow at the same rate.

An attractive feature of this framework is that its transition dynamics can be confronted directly with the data. This is precisely what Jones and Liu (2024) do in their Figure 4, where they back out the exact trajectories of β_t and Z_t from US time-series data on the labor share and output per hour from 1950 to 2020. One interpretation of these estimates is that they recover the technological progress in automation and quality improvements required for the model to fully account for the joint dynamics of output and the labor share. A natural next step, then, is to ask whether such technological progress can be microfounded as an equilibrium outcome of innovation and production decisions.

The answer is not obvious. Both Jones and Liu (2024) and the current paper provide a natural starting point by endogenizing automation and quality improvements through vertical and horizontal R&D investment decisions, following the tradition of Schumpeterian growth models. The two papers differ only in the R&D input: the former assumes that R&D uses the final good as lab equipment, whereas the latter assumes that it uses skilled labor. Let λ_t^v and λ_t^h denote the arrival rates of vertical and horizontal innovation, respectively. In both papers, the state variables β_t and s_t satisfy the following laws of motion:

$$\dot{\beta}_t = \lambda_t^h (1 - \beta_t), \qquad \dot{s}_t = \frac{\lambda_t^h}{q} - (\gamma - 1)\lambda_t^v s_t, \tag{1}$$

where q captures the productivity of newly automated tasks and $\gamma > 1$ is the step size of quality improvement.

At first glance, it may seem that one can choose λ_t^h and λ_t^v freely to match any desired trajectories of β_t and s_t. Building on the framework in the paper, however, my reading is that their joint evolution is not unconstrained. In equilibrium, λ_t^h and λ_t^v are pinned down by free-entry conditions and by the allocation of a common input, either final goods or skilled labor. As a result, the equilibrium values of λ_t^h and λ_t^v must equalize the marginal returns to vertical and horizontal innovation. In both models,

$$\frac{\lambda_t^h}{\lambda_t^v} = \left(\frac{\xi_h}{\xi_v}\right)^{\frac{1}{1-\theta}}, \tag{2}$$

where ξ_h (ξ_v) measures the efficiency of horizontal (vertical) innovation and $\theta < 1$ is the returns-to-scale parameter in R&D. Although both papers present this relationship only along the balanced-growth path, my extension of the framework suggests that it can be carried over to the entire transition path, since free entry and free factor mobility hold at every point in time.

This relationship places strong restrictions on aggregate labor-share dynamics. Using Equations (1) and (2), we obtain the steady state

$$s^* = \frac{1}{(\gamma-1)q}\left(\frac{\xi_h}{\xi_v}\right)^{\frac{1}{1-\theta}}, \tag{3}$$

which is independent of the level of λ_t^h (or λ_t^v). We can therefore rewrite the law of motion for s_t as

$$\dot{s}_t = (\gamma-1)\lambda_t^v(s^* - s_t).$$

It follows that if the initial state of the economy satisfies $s_0 = s^*$, as assumed in the paper's calibration, then s_t changes only when the steady-state value s^* changes. Because s^* is determined solely by the parameters of the innovation technology, this extension of the model implies that the aggregate labor share remains constant along the transition, even in the presence of business-cycle shocks or policy changes. Under this logic, the increase in skilled labor emphasized in the current paper would not be expected to affect the labor share.

Rather than viewing this strict constancy as a shortcoming, I think it is more helpful to interpret it as a benchmark implication that highlights the robustness of the cost-disease mechanism. Perhaps the reason we observe a stable labor share over long horizons is precisely that this mechanism is powerful enough to stabilize it quickly. An interesting next step, therefore, would be to quantify the strength of this mechanism. For example, one could relax free factor mobility by introducing adjustment costs, so that the ratio λ_t^h/λ_t^v is no longer constant, and then ask whether the model can generate the amount of labor-share fluctuation observed in the data.

In addition, this framework offers a useful perspective on the recent decline in the labor share across many countries: such a decline may reflect structural changes in innovation technology. Equation (3) shows that a lower labor share, equivalently a higher s^*, can arise either from a decline in the step size of quality improvement γ or from an increase in the relative efficiency of automation ξ_h/ξ_v. Both possibilities are consistent with the evidence in Bloom et al. (2020) that vertical innovation in advanced economies has become more difficult over the past few decades. Alternatively, global trade liberalization may have facilitated the diffusion of automation technologies, for example through the import of industrial robots, thereby increasing the relative efficiency of automation ξ_h/ξ_v. This trade-based explanation may be especially relevant for developing countries such as China in the early 2000s, where the labor share also declined rapidly (Chang et al., 2016). Future research could explore these channels in greater detail and evaluate their empirical relevance.

Data Availability Statement

Data sharing not applicable to this article as no datasets were generated or analysed during the current study

References

Acemoglu, D. and P. Restrepo, "The Race between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment," *American Economic Review* 108 (2018), 1488–1542.

Aghion, P., B. F. Jones and C. I. Jones, "Artificial Intelligence and Economic Growth," in *The Economics of Artificial Intelligence: An Agenda* (University of Chicago Press, 2019), 237–282.

Bloom, N., C. I. Jones, J. Van Reenen and M. Webb, "Are Ideas Getting Harder to Find?" *American Economic Review* 110 (2020), 1104–1144.

Chang, C., K. Chen, D. F. Waggoner and T. Zha, "Trends and Cycles in China's Macroeconomy," *NBER Macroeconomics Annual* 30 (2016), 1–84.

Jones, B. F. and X. Liu, "A Framework for Economic Growth with Capital-Embodied Technical Change," *American Economic Review* 114 (2024), 1448–1487.

Restrepo, P., "Automation: Theory, Evidence, and Outlook," *Annual Review of Economics* 16 (2024), 1–25.