A Report on "Distributional Growth Accounting: Education and the Reduction of Global Poverty, 1980–2019" by Gethin (2025)

Reviewer 2

February 13, 2026

v1

isitcredible.com

Disclaimer

This report was generated by large language models, overseen by a human editor. It represents the honest opinion of The Catalogue of Errors Ltd, but its accuracy should be verified by a qualified expert. Comments can be made here. Any errors in the report will be corrected in future revisions.

> I am wiser than this person; for it is likely that neither of us knows anything fine and good, but he thinks he knows something when he does not know it, whereas I, just as I do not know, do not think I know, either. I seem, then, to be wiser than him in this small way, at least: that what I do not know, I do not think I know, either.
>
> Plato, *The Apology of Socrates*, 21d

> To err is human. All human knowledge is fallible and therefore uncertain. It follows that we must distinguish sharply between truth and certainty. That to err is human means not only that we must constantly struggle against error, but also that, even when we have taken the greatest care, we cannot be completely certain that we have not made a mistake.
>
> Karl Popper, 'Knowledge and the Shaping of Reality'

Overview

Citation: Gethin, A. (2025). Distributional Growth Accounting: Education and the Reduction of Global Poverty, 1980–2019. *Quarterly Journal of Economics*, Vol. 140, No. 4, pp. 2571–2618.

URL: https://doi.org/10.1093/qje/qjaf033

Abstract Summary: This article quantifies the role played by education in the reduction of global poverty from 1980 to 2019 using a new distributional growth accounting framework and a comprehensive microdatabase.

Key Methodology: Distributional growth accounting framework combining standard growth accounting with a model of education and the wage structure (\u00e0 la Goldin and Katz, 2007), exploiting a new microdatabase representative of nearly all of the world's population, new estimates of the private returns to schooling, and historical income distribution statistics. The methodology is validated using quasi-experimental evidence from three large-scale schooling initiatives.

Research Question: What role has education played in the reduction of global poverty between 1980 and 2019?

Summary

Is It Credible?

This study presents a formidable effort to quantify the economic impact of the global expansion of education over the last four decades. By assembling a new micro-database covering 97% of the world's population and developing a "distributional growth accounting" framework, Gethin challenges the consensus derived from standard macroeconomic models. The article's headline claim is striking: "Education can account for about 45% of global economic growth and 60% of pretax income growth among the world's poorest 20% from 1980 to 2019" (p. 2571). This asserts that human capital accumulation has been the primary engine of poverty reduction, far exceeding the 16% contribution estimated by standard methods (p. 2598). The analysis argues that previous approaches failed because they ignored how an increasing supply of skilled labor reduces the skill premium, thereby disproportionately benefiting low-income, low-skilled workers.

The credibility of the qualitative argument—that standard growth accounting underestimates the benefits of education for the poor—appears high. The mechanism of imperfect substitution, where an increase in the supply of educated workers compresses the wage structure, is economically sound and supported by the model's logic. By accounting for within-country inequality and the fact that the poor rely almost entirely on labor income, the framework plausibly demonstrates that the "standard" 16% estimate is a lower bound.

However, the precision and magnitude of the headline quantitative claims—specifically the "60%" figure for the global poor—are subject to significant uncertainty. The analysis relies on a benchmark calibration for the elasticity of substitution between skill groups (σ), a parameter that is notoriously difficult to pin down. The article acknowledges this sensitivity: varying the elasticity parameters within a plausible range derived from the literature shifts the estimated contribution of education to the growth of the poor from 47% to over 100% (p. 2602). Consequently, the specific point estimate of 58% (rounded to 60% in the abstract) represents one possibility within a wide distribution of potential outcomes.

Furthermore, the foundational assumption that educational expansion is exogenous to economic growth (p. 2580) likely introduces an upward bias to these estimates. If economic growth, state capacity, or technological change drive the expansion of schooling, attributing the subsequent income gains solely to education conflates cause and effect. While the article is transparent about this limitation, noting that "educational choices themselves may be shaped by skill-biased technical change" (p. 2579), the magnitude of this reverse causality remains unquantified. If a significant portion of the schooling boom was a response to growth rather than its cause, the causal contribution of education would be lower than the accounting contribution presented.

The robustness of the specific poverty reduction estimates also depends on data selection. When analyzing the decline in extreme poverty (Table IV), the article utilizes World Bank data, which yields a result where education explains 67% of the decline at the $6.85 threshold (p. 2597). However, using the World Inequality Database—which the author describes elsewhere as "a more adequate source" for this type of analysis (Supplementary Online Appendix, p. 75)—yields a considerably lower estimate of 44% (Supplementary Online Appendix, p. 5). The presentation highlights the more favorable estimate in the main text, which emphasizes the upper bound of education's impact.

Finally, the narrative regarding skill-biased technical change (SBTC) requires nuance. The abstract states that a "significant fraction of these gains was made possible by skill-biased technical change amplifying the returns to education" (p. 2571). While true for the global average (driven by high-income countries), the analysis finds that for the global poor—concentrated in China and India—the interaction between education and technology was actually negative during the 2000–2019 period (p. 2611). For these populations, education reduced poverty *despite* labor demand trends, not because of them.

In summary, the article successfully establishes that the distributional benefits of education are substantial and likely larger than previously understood. The claim that education is a key driver of poverty reduction is credible. However, the specific quantification that it accounts for 60% of income growth for the poor should be viewed as a likely upper-range estimate, contingent on specific parameter choices, data sources, and the assumption of exogeneity.

The Bottom Line

Gethin provides a compelling methodological advance that likely corrects a long-standing underestimation of education's role in global poverty reduction. The central finding—that increasing the supply of education compresses wage inequality and disproportionately benefits the poor—is economically robust and well-supported by the framework. However, the headline figure that education explains 60% of growth for the poorest 20% is highly sensitive to modeling assumptions and data selection, and it relies on the strong premise that educational expansion is independent of economic growth. Readers should interpret the results as strong evidence that education is a major driver of development, while treating the precise point estimates with caution.

Potential Issues

Exogenous educational expansion: The study's central counterfactual exercise rests on the foundational assumption that the global expansion of education from 1980 to 2019 was an exogenous event, independent of the economic growth it is purported to explain. The author explicitly states this choice: "Throughout the article, I treat education as exogenous and leave the study of the determinants of schooling for future research" (p. 2580). This assumption is debatable, as economic growth, technological change, and increased state capacity are themselves powerful drivers of both the demand for and supply of education. If the same underlying forces that drive income growth also drive educational expansion, the model may attribute a portion of that growth to education due to reverse causality, potentially overestimating education's true causal effect. The author acknowledges this limitation, noting that "educational choices themselves may be shaped by skill-biased technical change" but argues that quantifying this channel is not possible without historical survey data from the 1980s (p. 2579). While this transparency is commendable, the unmodeled endogeneity remains a significant potential issue that could affect the magnitude of the article's headline quantitative claims.

Sensitivity of results to elasticity of substitution: The article's main quantitative findings, particularly the distributional effects for the poor, are highly sensitive to the calibrated values for the elasticity of substitution (σ) between different skill groups. The benchmark specification ($\sigma_1 = 4$, $\sigma_2 = 6$, $\sigma_3 = 8$) yields the headline result that education explains 58% of income growth for the world's poorest 20% (p. 2596). However, the author's own sensitivity analysis reveals that this estimate is fragile. As shown in Table VI (p. 2602), other plausible parameter choices, also drawn from the literature, produce a wide range of results for this group: from 47% under a "high substitutability" scenario to 109% under a "very low substitutability" scenario. The upper bound of this range (109%) corresponds to a specific scenario using shortrun elasticities, which assumes no endogenous adjustment in labor demand. This wide range implies that reasonable changes in this uncertain parameter can alter the conclusion from education explaining less than half of income growth for the poor to it explaining more than all of it. While the author is transparent in providing this sensitivity analysis and justifies the benchmark choice by citing recent literature (p. 2591), the presentation of a precise "about... 60%" figure in the abstract (p. 2571) may mask the significant uncertainty underlying this point estimate.

Sensitivity to benchmark modeling assumptions: The article's headline estimates are sensitive to two key modeling choices that are presented as "conservative" but are debatable and have large effects on the results. First, the benchmark model assumes that physical capital is not affected by schooling (p. 2593). A robustness check that allows capital to adjust to maintain a constant capital-output ratio, a specification described as "standard in the literature" (p. 2604), increases the estimated contribution of education to global growth from 45% to 62%. Second, the benchmark assumes that returns to schooling apply to 100% of mixed income. A sensitivity analysis shows that assuming returns apply to only 75% of mixed income—an approach the author notes is "common in the literature" (p. 2600, footnote 17)—would reduce the estimated contribution for the poorest 20% from 58% to 46% (Supplementary Online Appendix, p. 6). The decision to select a benchmark that includes these specific assumptions significantly shapes the final quantitative claims.

Data construction and imputation: The construction of the novel global microdatabase, a core contribution of the article, involves several methodological choices and imputations that introduce unquantified uncertainty. For six countries where individual income data was unavailable, income was proxied by "splitting equally household expenditure among adults in employment," a method that assumes no intra-household inequality and that expenditure is a good proxy for pre-tax income (Supplementary Online Appendix, p. 56). Furthermore, for a substantial number of countries (e.g., 66 countries for primary education returns), returns to schooling could not be estimated directly and were imputed based on the average ratio of returns across education levels in other countries (Supplementary Online Appendix, p. 63). The author is transparent about these procedures and the variable quality of the underlying surveys (p. 2607). However, the uncertainty stemming from these necessary but ad-hoc harmonization and imputation choices is not formally incorporated into the final estimates, making their precision difficult to assess.

Narrative on skill-biased technical change: The abstract and introduction propose a key narrative that "A significant fraction of these gains was made possible by skill-biased technical change amplifying the returns to education" (p. 2571). This framing reflects the finding that, for the average country, about 30% of the benefits of education were enabled by skill-biased technical change (p. 2575). However, this global average masks important heterogeneity. The article's own results for the 2000–2019 period show the opposite for China and India, the two countries most central to the global poverty reduction story. For these countries, the analysis finds that the interaction between education and technology was negative, meaning "education would have had larger effects if labor demand had remained constant" (p. 2611). While the author discusses this nuance in the text, the abstract's summary of the average effect may not fully capture the mechanism at play for the majority of the population that escaped extreme poverty during this period.

Inconsistent use of data sources for poverty analysis: The article uses two different global income datasets for its main analyses, leading to different conclusions about poverty reduction. The primary distributional growth analysis relies on the World Inequality Database (WID), which the author states is "a more adequate source" for the study's purposes (Supplementary Online Appendix, p. 75). However, for the main poverty reduction analysis in Table IV, the article switches to World Bank data, justified on the grounds that it is the "most commonly used data source to measure extreme poverty" (p. 2597). This switch is consequential: at the $6.85/day poverty line, the World Bank data implies that education explains 67% of poverty reduction, whereas the WID data implies it explains only 44% (pp. 2597; Supplementary Online Appendix, p. 5). The article features the substantially larger estimate derived from World Bank data in the main text, while the smaller estimate derived from the author's preferred dataset is relegated to an appendix. This presentation choice emphasizes the upper bound of the estimated effect.

Model validation: The article presents an exercise to validate its framework against three quasi-experimental studies of education expansion (pp. 2593–2595), but the strength of this validation may be overstated. For each case, the model is calibrated using the aggregate return to schooling estimated from the specific context it is meant to validate (p. 2595), making the exercise more of a check on the model's internal distributional mechanics rather than an independent test of its predictive power. Indeed, the model's ability to replicate aggregate effects is circular by construction. Furthermore, in the U.S. case, the model significantly underestimates the observed causal effect for low-income groups, a finding the author attributes to human capital externalities not included in the model (Supplementary Online Appendix, p. 22). While the exercise provides some support for the model's core mechanism—that supply effects benefit the poor—the article's description of the model performing "remarkably well" (p. 2594) may not fully capture these nuances.

Omission of signaling effects: The framework assumes that private returns to schooling reflect true increases in productivity-enhancing human capital. The author briefly acknowledges that the analysis could "overestimate them in the presence of signaling effects" (p. 2593), where education serves primarily to reveal pre-existing ability rather than create it. If signaling is a major component of the returns to education, the social return and contribution to aggregate growth would be lower than the private returns suggest. While the article's validation exercise, which examines aggregate regional income, provides some indirect evidence against a pure signaling story, the model does not formally account for this standard alternative explanation, which remains an acknowledged limitation.

Presentation and clerical issues: Several minor presentation and clerical issues appear in the article. First, the analysis of skill-biased technical change (Section V) relies on a separate, smaller, and lower-quality dataset, a fact that is noted in the text (p. 2607) but not in the abstract, where the finding is presented alongside the main results. Second, there is a minor rounding difference between the abstract and the main results table: the abstract states that education explains "about... 60% of pretax income growth" for the poorest 20% (p. 2571), while the main results table reports the figure as 58% (p. 2596). Third, several tables contain apparent calculation discrepancies because the displayed numbers are rounded. For example, in Table III (p. 2596), the reported contribution (1.1) and growth (1.5) for the "Middle 40%" would imply a share of 73.3%, but the table reports 69%. This is likely an artifact of the "Share" column being calculated from unrounded underlying data while the other columns are rounded for display, but it can create confusion for the reader.

Future Research

Endogenizing educational expansion: Future work could integrate a model of educational demand to address the reverse causality between income growth and schooling. By explicitly modeling how economic growth and state capacity enable educational expansion, researchers could isolate the distinct causal contribution of schooling to poverty reduction, separating it from the broader development process that drives both.

Incorporating signaling effects: Future studies could refine the growth accounting framework by incorporating signaling models of education. If a portion of the private return to schooling reflects ability revelation rather than human capital accumulation, the social return to education would be lower than the private return. Quantifying this wedge would provide a more accurate estimate of education's contribution to aggregate economic output.

Probabilistic parameter estimation: Future research could employ Bayesian methods to formally incorporate the uncertainty surrounding key parameters, such as the elasticity of substitution and the returns to schooling. Rather than relying on a single benchmark calibration with sensitivity checks, this approach would generate a probability distribution of the contribution of education to growth, providing policymakers with a clearer understanding of the range of likely effects.

© 2026 The Catalogue of Errors Ltd

This work is licensed under a

Creative Commons Attribution 4.0 International License

(CC BY 4.0)

You are free to share and adapt this material for any purpose,
provided you give appropriate attribution.

isitcredible.com