Globalisation, Income, and Welfare: Imported Goods Variety and Quality*

Céline Azémar[a], Rodolphe Desbordes[b], and
Markus Eberhardt[c]

[a]Rennes School of Business
[b]SKEMA Business School, Paris and Université Côte d'Azur, Nice
[c]University of Nottingham and Centre for Economic Policy Research

This version: July 10, 2026

Abstract: Did globalisation benefit consumers in poor countries? Using granular data on trade in goods for 138 countries, we measure the welfare gains from wider variety and higher quality of imported goods over the past quarter-century. The variety gains are large, a median 19 percent fall in the cost of living, and strongly pro-poor: poorer countries gained twice as much as richer ones. Because they derive from falling trade costs, variety gains made poor countries richer in real consumption even in the absence of income growth. Quality, by contrast, tracks income, and because income barely converged, neither did quality.

Keywords: gains from trade, import variety, globalisation, convergence, consumption

JEL codes: F14, F61, O19, D63

*In preparing this paper, the authors used an AI-based assistant tool (Anthropic's Claude) to support coding, the construction of some of the data, the generation of tables and figures, and language editing. The research design, the empirical analysis and its specification, the drafting and writing of the manuscript, the interpretation of the results, and all conclusions are the authors' own. The usual disclaimers apply. Correspondence: Markus Eberhardt, Sir Clive Granger Building, University of Nottingham, University Park, Nottingham NG7 2RD, UK; email: markus.eberhardt@nottingham.ac.uk

1 Introduction

"Foreign trade. . . [is] highly beneficial to a country, as it increases the amount and variety of the objects on which revenue may be expended."
David Ricardo, *On the Principles of Political Economy and Taxation* (1817, Ch. 7)

The political economy of globalisation has been written almost entirely on the production side: import competition displaced manufacturing workers (the 'China shock'; Autor et al. 2013, 2016, Pierce and Schott 2016), and the gains and losses from trade fell unevenly across regions, sectors, and skills (Helpman et al. 2016, Dorn and Levell 2024, Rodrik 2024). This is the globalisation that produced a popular backlash, studied extensively in its socio-economic and political implications. A separate literature has quantified the effects of globalisation, through increased exports and more export destinations, on exporters and the countries in which they are located (Melitz 2003, Eaton et al 2011, Brambilla et al. 2012, De Loecker 2013, Arkolakis et al. 2021). But, as David Ricardo noted over 200 years ago, trade also reaches people as *consumers*, through lower prices and access to a wider range and/or quality of goods. On this consumer margin, Broda and Weinstein (2006) estimated that the United States gained on the order of 3 percent of GDP over three decades (1970-2001) from a wider import variety. Their estimate, and most of what we know about consumer gains, applies to rich economies (Feenstra 2018, Lo Forte 2025).[1]

What did the consumer side of globalisation look like for the rest of the world? Which countries gained from increased variety and quality? These questions represent the between-country counterpart of the within-country 'China shock' debate. Their answers could substantially change the progressivity we assign to the distributional record of trade and globalisation in contemporary and historical periods: if consumer benefits accrued mainly to already-rich economies, global integration widened the welfare gap between nations; if benefits disproportionally reached poor countries, the gap narrowed. In this paper, we measure the welfare gains from imported variety and quality using BACI HS6 data (around 5,000 product categories) for 138 countries over 1995–2023 and ask how the respective gains are distributed between richer and poorer economies. Additional analysis uses less granular SITC-4 data (around 1,000 categories) but covers an extended time horizon (1962–2023).

We find that the variety gains are large and strongly progressive. Applying the Feenstra (1994) exact price index, with substitution elasticities estimated directly product by product from global trade flows, we find a median welfare gain worth about 19 percent of the cost of living, with poorer countries gaining about twice as much as richer ones.[2] The pro-poor pattern is robust to restricting the analysis to consumer goods or to manufactures (in both cases welfare gains *increase*), excluding entrepot exporters, and reweighting the products. Conditioning on initial variety identifies the mechanism: countries that began being variety-

[1] An exception is Arkolakis et al. (2008), who find only modest variety gains for Costa Rica.

[2] These are our preferred, conservative figures, computed after dropping small trade flows which most likely to reflect changes in reporting coverage; before those flows are removed the pro-poor incidence is steeper still.

poor caught up over time, and the originally variety-deprived were disproportionately poor. This is convergence in variety *access*, the extensive-margin analogue of income convergence.

This gain from variety accrues to real consumption without requiring income growth. It is a fall in the cost of living driven by falling trade costs, not by a country's own income, and conventional income statistics, built on price indices that miss new varieties, do not record it. The point is most pronounced once we exclude China and India, which had strong income growth: for the remaining 136 economies, between-country incomes barely converged at all over 1995–2023, yet the variety channel narrowed real-consumption gaps as much as ever. Poor countries became richer in what they consume without becoming richer in income. In our extension we use SITC-4 data from The Atlas of Economic Complexity and study ten-year windows between 1962 and 2023: we find that pro-poor convergence in variety access is concentrated in the 1990s and 2000s, dropping to virtually zero in the most recent period.

Did this convergence process also extend to quality? Variety is the extensive margin; quality, whether one buys the dearer or the cheaper version of each good, is the intensive one. We measure it with ϕ, a trade-revealed index of the consumption quality tier built from the within-product unit values of everyday-consumption imports (see Section 3). The contrast is sharp: across a range of alternative variants of the consumption quality index, we find that quality did *not* converge for poorer countries, because it follows income and (excepting India and China) income barely converged during our time period. Changes in the quality levels of imports did occur, but these were not conducive to bringing welfare gains to poorer countries. On average, poor countries experienced more choice, but not more quality.

The paper makes three contributions. First, it takes the gains-from-variety accounting of Feenstra (1994) and Broda and Weinstein (2006), applied mostly to individual economies, to the full cross-country distribution. The closest precedent, Fajgelbaum and Khandelwal (2016), measures unequal gains *across consumers within* a country; we measure them *across* countries. Second, it pairs the variety margin with a trade-revealed quality tier and traces their contrasting fates to a simple non-homothetic framework: access follows trade costs, quality follows income. Third, it reads variety access as an extensive-margin counterpart to income β-convergence (Patel et al. 2021, Kremer et al. 2022, Johnson and Papageorgiou 2020). This channel alone reduced between-country inequality, a progressive force that the production-side account of the globalisation backlash overlooks. The paper further builds on the measurement of heterogeneous price indices (Hottman et al. 2016), the consumer gains from integration in developing countries (Atkin et al. 2018), the extensive margin of trade (Hummels and Klenow 2005), and the welfare-gains-from-trade framework (Arkolakis et al. 2012). We do not claim to conduct the first cross-country variety measurement; however, the distributional incidence and the variety-versus-quality contrast are, to our knowledge, new.

The remainder of this paper is organised as follows: Section 2 sets out the framework, Section 3 the data and indices; Sections 4–5 establish the variety gains and contrast them with the quality margin. Section 6 concludes.

2 Two margins, two drivers

A simple framework organises the two margins and shows why they should respond to different forces. Consider a consumer in importer j with real expenditure y_j who buys differentiated varieties of each good. A variety from exporter e has quality z_e and reaches j at the delivered price $\tau_{je}\, p_e$, where $\tau_{je} \geq 1$ is an iceberg trade cost. Globalisation is a fall in τ.

The extensive margin: access. A variety is consumed only if it is worth importing. With a fixed cost of sourcing each variety, e enters j's basket only when its delivered price is low enough to cover that cost, that is, when it clears a reservation level,

$$\tau_{je}\, p_e \leq \bar{p}_{jg}, \tag{1}$$

so the set of accessible varieties $N_{jg}(\tau)$ widens as τ falls and is bounded above by the world set of producers. Two things follow. Lower trade costs raise access, $\partial N/\partial \tau < 0$; and because access is bounded, economies that began with high trade costs, and so few varieties, gain proportionally the most as τ falls, so N_{jg} converges across countries as trade costs converge. What determines access is τ, not income: at a common trade cost a poor and a rich country reach almost the same set, with income entering only through second-order market-size effects. The Feenstra variety term of Section 3 measures the welfare value of this expanding variety of goods.

The intensive margin: quality. Within the accessible set the consumer chooses a quality tier. Let preferences be non-homothetic in quality, so that the marginal valuation of quality rises with income and the chosen tier q^*_{jg} satisfies

$$\frac{\partial q^*_{jg}}{\partial y_j} > 0, \qquad \frac{\partial q^*_{jg}}{\partial \tau_{je}} \approx 0. \tag{2}$$

Quality is a superior characteristic, so the tier rises with income; and because an iceberg cost scales the price of every quality equally, it leaves the quality choice essentially unchanged. The index ϕ of Section 3 is the empirical counterpart of quality tier q^*.

The prediction. A fall in τ raises access for everyone and equalises it, while leaving the quality tier where income puts it. Convergence in access without convergence in quality is therefore the diagnostic pattern of a trade-cost decline under non-homothetic preferences: the extensive margin follows trade costs, the intensive margin follows income. The rest of the paper measures the two margins and finds exactly this. Non-homothetic demand for quality is the mechanism in Hallak (2006), Fieler (2011) and Fajgelbaum and Khandelwal (2016); the extensive-margin response to trade costs is Hummels and Klenow (2005) and Feenstra (1994).

3 Data and measurement

Trade data. The data for our main results are taken from BACI (Gaulier and Zignago 2010), CEPII's reconciliation of UN Comtrade bilateral flows, in HS92 over the 1995 to 2023 period (around 5,000 product categories). For the purposes of our study, a *variety*, in the sense of Broda and Weinstein (2006), is a country-product pair: a given HS6 good supplied by a given exporter. For each importer we observe the value and quantity of each variety it buys in each year. In our baseline results, we take three-year windows at each end, 1995 to 1997 and 2021 to 2023, to dampen year-to-year noise in variety entry and exit.

In an extension expanding the time horizon of our analysis at the cost of less granular product categories, we use SITC-4 (around 1,000 product categories) bilateral trade flows from The Atlas of Complexity and study ten-year windows over 1962-2023.

Income data. For the income gradient of goods variety and quality, and the 'poor'/'rich' sample splits, we use PPP GDP per capita from the World Bank's World Development Indicators; income growth over 1995–2023 is the log change in this series between the two three-year windows. The historical extension (Section 4.3) instead uses expenditure-side real GDP per capita (variable *rgdpe* per capita) from the Penn World Table (version 10.0), which is more widely available back to 1962.

The variety price index. A conventional price index measures the cost of a *fixed* basket and is, by construction, blind to new varieties: a good, or a new source for a good, that a country did not import in the base period, has no base-period price, so a fixed-basket index registers nothing when it first appears, even though gaining access to it makes consumers better off. Feenstra (1994) addresses this issue by deriving the exact cost-of-living index for constant-elasticity-of-substitution (CES) preferences when the set of available varieties itself changes between periods.[3] For a given importer and good g with elasticity of substitution $\sigma_g > 1$, the exact price index between the base period 0 and the end period 1 factors into two pieces,

$$\frac{P_{g,1}}{P_{g,0}} = \underbrace{\prod_{e \in C_g} \left(\frac{p_{ge,1}}{p_{ge,0}} \right)^{\omega_{ge}}}_{\text{conventional index, common varieties}} \times \underbrace{\left(\frac{\lambda_{g,1}}{\lambda_{g,0}} \right)^{1/(\sigma_g - 1)}}_{\text{variety adjustment}}, \tag{3}$$

where C_g is the set of varieties (here, exporters) supplying good g in *both* periods, the ω_{ge} are ideal log-change (Sato–Vartia) weights that sum to one over C_g,[4] and $\lambda_{g,t}$ is the share of period-t expenditure on good g that falls on the common set C_g. The first factor is an ordinary price index computed only over the varieties one can actually compare across periods; the second is the correction for the varieties that entered or exited.

[3] This is a cost-of-living index isolating the consumer-variety channel from observed import shares, not a gains-from-trade-against-autarky measure in the manner of Arkolakis et al. (2012).

[4] Sato–Vartia weights are logarithmic means of each variety's expenditure share across the two periods, chosen so that the resulting index satisfies standard index-number consistency properties.

All of the economics is in the second factor, and its appeal is that it needs no information on the new varieties themselves. A complete welfare calculation would seemingly need each new variety's reservation (virtual) price, which under CES is formally infinite and never observed; yet this is not needed. The expenditure share retained by the surviving common varieties is a sufficient statistic for the welfare value of everything that entered or exited: When entrants capture much of end-period spending, the common share $\lambda_{g,1}$ falls well below $\lambda_{g,0}$ and the index infers the new varieties were valuable. Summing the negative log of the variety factor across goods, weighted by current-period expenditure shares w_{jg}, gives importer j's welfare gain from the expansion of variety,

$$\ln G_j = \sum_g w_{jg} \frac{1}{\sigma_g - 1} \ln \frac{\lambda_{jg,0}}{\lambda_{jg,1}}, \tag{4}$$

in which $\lambda_{jg,t}$ now carries the importer subscript j. New entry lowers the end-period common share $\lambda_{jg,1}$ and so raises $\ln G_j$, and the gain is larger the lower is σ_g: the exponent $1/(\sigma_g - 1)$ is largest exactly when varieties are poor substitutes, so that an extra one is worth most, and it vanishes as $\sigma_g \to \infty$ and the varieties become homogeneous.

Estimating the elasticities. The welfare gain depends on the elasticity of substitution σ_g, so rather than calibrate a single value we estimate it for each HS6 good by the Feenstra (1994) method applied to world trade, treating exporters as the competing varieties. This estimation faces a standard identification problem: within a good, consumer substitution across exporters traces out a downward-sloping CES demand curve while exporters' costs trace out an upward-sloping supply curve, and the data lie only at their intersection, so a single market cannot separate the two. Feenstra's insight is that a panel of varieties can: differencing each variety against a reference exporter leaves variety-specific demand and supply shocks, and, assuming these are orthogonal, the cross-variety co-movement of price and share dispersion identifies both elasticities. That heteroskedasticity across varieties is the identifying variation. We solve the resulting moment conditions good by good, with a grid search over the structural parameters where the closed-form solution is infeasible, following Broda and Weinstein (2006). The procedure yields elasticities for 4,970 HS6 goods with a median of 3.6, and 97 percent between 2 and 10. The central tendency is consistent with the Broda and Weinstein (2006) distribution (US median near 3). The estimated σ_g enter (4) to compute importer welfare gains.

The quality tier. The intensive margin asks whether a country buys the dearer or the cheaper version of a given good. We measure it with ϕ, a trade-revealed quality index built from within-product unit-value comparisons. Let $u_{egj,t} = v_{egj,t}/q_{egj,t}$ be the unit value (value divided by quantity) of the HS6 good g that importer j buys from exporter e in year t.[5] These

[5]BACI values are reported free on board: CEPII estimates and removes the CIF freight-and-insurance margin when reconciling importer and exporter declarations (Gaulier and Zignago 2010). Unit values are

unit values are computed on a fixed basket of everyday final-consumption goods, defined by mechanical filters on the HS6 lines: the basket keeps lines that are (i) final-consumption non-durables (BEC categories 112, 62 and 63), (ii) high value-to-weight (a pooled median unit value of at least \$5,000 per tonne, so that a recorded unit value is informative about the good rather than dominated by weight), and (iii) broadly traded (imported by at least 60 percent of countries in at least 80 percent of years), leaving a few hundred consumer goods whose within-product unit values are comparable across importers. Separately in each year, ϕ is the importer effect in the weighted two-way fixed-effects regression

$$\ln u_{egj,t} = \delta_{eg,t} + \phi_{jt} + \varepsilon_{egj,t}, \tag{5}$$

estimated by weighted least squares with weights equal to the trade value $v_{egj,t}$, winsorised at the 99th percentile so that a few very large flows do not dominate. The work is done by the exporter-by-product fixed effect $\delta_{eg,t}$: it absorbs the price of good g as supplied by exporter e, so that ϕ_{jt} is identified purely from *within-cell* variation, i.e., from how much importer j pays for the very same good from the very same source relative to other importers. Differences in product mix across countries and in suppliers' comparative-advantage prices are thereby differenced out; what remains in ϕ_{jt} is the quality tier the importer selects within each narrow good, because an HS6 line from a single exporter still bundles cheaper and dearer variants and a higher unit value picks out the dearer, higher-specification (quality) one.

The regression is run on importers with at least 200 distinct exporter-product cells, so that the importer effect is precisely estimated. Because a two-way fixed-effects model identifies ϕ only up to an additive constant, the index is normalised to a trusted frontier: the mean of ϕ over a set of high-income reference countries is set to zero, and each country's ϕ_{jt} is read relative to that frontier. A higher ϕ therefore means a country systematically buys the dearer, higher-quality varieties of the same goods from the same suppliers, and the change in ϕ between 1995 and 2023 is its movement up or down the consumption quality ladder. Being a unit-value object, ϕ is noisier than the variety counts; Section 5 shows the quality result is robust across five constructions of ϕ that address this noise, including a size-adjusted version that adds a log-quantity covariate, a three-year-pooled estimate and a Kalman-smoothed series.

therefore net of freight, so ϕ does not mechanically capture the higher shipping cost a remote importer pays for the same good; and because the CIF/FOB conversion is distance- and product-specific, only second-order residual freight remains. The within-cell design nets out any freight common to a supplier-product in any case.

4 The gains from variety

4.1 The pro-poor incidence

Throughout, the income *gradient* of an outcome is its cross-country slope on initial log PPP income per capita (estimated by ordinary least squares, with the simple correlation reported alongside); a negative gradient is *pro-poor*, meaning poorer countries gain more. This differs from a *margin*, which names the dimension of consumption at issue: the extensive margin of variety, or the intensive margin of quality.

Main results. Panel A of Table 1 reports the variety results on our preferred ('conservative') specification, which drops the small, reporting-sensitive flows. The median country's cost of living fell by about 19 percent through the expansion of its import variety over 1995 to 2023 (about 15 percent before those flows are removed), several times the rich-country magnitude in Broda and Weinstein (2006) because the developing world integrated from a far narrower starting point. For the rich economies they studied, our estimates are of a similar order of magnitude: the United States gains 2.5 percent, almost exactly the 2.6 percent Broda and Weinstein report for the United States over 1972 to 2001, and Japan 3.5 percent. Developing economies in our sample gain far more: China and India both around 13 percent.

The welfare gains fall steeply with income. The log gain regressed on initial log income has a slope of -0.053 (correlation -0.49).[6] The top-left panel of Figure 1 plots the welfare gains against income, the top-right panel the variety counts. The poorer half of countries gained 24.7 percent against 11.7 percent for the richer half, over double the effect size. Importantly, the benefits are near-uniform among the poorer countries, given that the income gradient turns insignificant, -0.013; they still display statistically significant progressivity in the richer half of the sample, -0.061.

Heterogeneity and Robustness. Panel B of Table 1 confirms that these patterns do not depend on any single modelling choice. Restricting the sample to consumer goods, the cleanest object to measure consumer-gains, gives -0.058; excluding the major entrepot exporters, -0.055; weighting products equally, a steeper -0.068; and a 2000 to 2019 window (to focus on the most recent time period while isolating the estimate from a potential Covid bias) a flatter but still clearly progressive -0.045.

When we further separate out necessities from luxury-type goods the gradient for the former is over twice the magnitude, -0.092, relative to the latter, -0.044: our results are not clearly driven by the economic elite in poorer countries.

[6] The gradient's standard error is robust to the goods dimension. Block-bootstrapping it over the roughly 4,600 HS6 goods that enter the gain (1,000 resamples, the same draw applied to every importer) gives a standard error of 0.009, close to the analytic 0.008, and a 95 percent interval of $[-0.072, -0.039]$. The pro-poor gradient is not carried by a handful of goods, nor does it rest on treating the estimated σ_g as known.

Table 1: Welfare gains from import variety: incidence and robustness (1995–2023)

Panel A. Welfare gain by income

	Median gain (% cost of living)	Gradient on initial ln income (s.e.)	Correlation
All countries	19.4	-0.053^{***} (0.008)	-0.49
Rich half	11.7	-0.061^{***} (0.020)	-0.36
Poor half	24.7	-0.013 (0.027)	-0.07

Panel B. Robustness of the pro-poor income gradient

Variant	Median gain (%)	Gradient (s.e.)	Correlation
All goods (main)	19.4	-0.053^{***} (0.008)	-0.49
Consumer goods only	20.2	-0.058^{***} (0.012)	-0.37
Excluding entrepots	20.6	-0.055^{***} (0.007)	-0.52
Products weighted equally	22.6	-0.068^{***} (0.007)	-0.61
Window 2000–2019	10.9	-0.045^{***} (0.007)	-0.35
Consumer goods, necessity-type	18.6	-0.092^{***} (0.012)	-0.59
Consumer goods, luxury-type	16.8	-0.044^{**} (0.019)	-0.19
Calibrated $\sigma = 3$	30.5	-0.099^{***} (0.016)	-0.52
Calibrated $\sigma = 5$	15.3	-0.049^{***} (0.008)	-0.52
Calibrated $\sigma = 8$	8.7	-0.028^{***} (0.005)	-0.52
HS Chapters 25-97	20.1	-0.063^{***} (0.010)	-0.49
HS Chapters 28-97	24.1	-0.081^{***} (0.011)	-0.53
HS Chapters 84-90	25.9	-0.092^{***} (0.012)	-0.54
All flows (upper bound)	15.1	-0.084^{***} (0.012)	-0.55
Drop flows below \$10k	15.1	-0.061^{***} (0.009)	-0.53
Drop flows below \$250k	21.6	-0.042^{***} (0.009)	-0.37
Conditional on ln population		-0.054^{***} (0.008)	

Notes: 138 importers with population above one million; the gain is the Feenstra variety term (4) with estimated σ_g, a cost-of-living reduction over 1995–2023 on the conservative flow set that drops bilateral flows below \$50,000. *Panel A*: rich/poor split at the median of initial income, each half's statistics estimated within that half; the all-flow figures (median 15.1, gradient -0.084) are an upper bound, shown in Panel B. *Panel B*: every variant is on the conservative set except the final block, which shows sensitivity to the flow threshold; the necessity/luxury rows split consumer goods at the median world unit value; the calibrated-σ rows leave the incidence unchanged (correlation -0.52 throughout, against -0.49 at the estimated σ_g); the final row controls for log population. Standard errors are heteroskedasticity-robust; *, ** and *** denote significance at the 10, 5 and 1 percent level.

A natural concern is that the convergence in variety could instead be carried by agricultural and primary goods: poor countries may simply be importing more kinds of fruit. Our results show that the opposite is true. Restricting to HS chapters 25 to 97 (manufactures, minerals and fuels) results in a steeper gradient of -0.063, restricting to HS chapters 28 to 97 (chemicals onward, dropping minerals and fuels) -0.081, and restricting to HS chapters 84 to 90 (machines, transport equipment and instruments) -0.092. The progressive variety gradient is, if anything, *stronger* in manufactures than in the full sample.

The income progressivity finding does not hinge on the estimated elasticities. A common calibrated σ rescales every country's gain by the same factor $1/(\sigma - 1)$, so it moves the magnitude of the gain and its slope mechanically, but leaves the incidence unchanged: the correlation between the variety gain and initial income is -0.52 for calibrated values of 3, 5 and 8 alike, against -0.49 for the estimated σ_g (see Panel B of Table 1, rows for $\sigma = 3, 5, 8$).

Another concern is that BACI reconciles and drops very small flows, so a poor country's imports of a good that sat below the reporting floor in 1995-97 and rose above it by 2021-23 would enter as a spurious new variety, mechanically inflating poor-country gains. For this reason our main specification drops flows below \$50,000. Removing varieties worth less than \$10,000, \$50,000 and \$250,000 gives gradients of -0.061, -0.053 and -0.042, against -0.084 on all flows; the strong progressive incidence remains intact.

Finally, these estimates apportion equal weight to each country, so that they may be distorted by an effect dominating in small economies. We address this concern by including the log of initial population as an additional control in the main regression, resulting in a virtually identical estimate, -0.054.

Mechanism. What drives the pro-poor incidence is the starting point. To see this, we regress the growth in a country's *number* of imported varieties on its initial log variety count: the slope is -0.386 (correlation -0.74),[7] so variety *levels* converged strongly, and economies that began variety-poor added proportionally far more over the sample period. Initial variety was in turn tightly linked to income, the two correlated 0.77, because protected, autarkic, or simply poor economies imported from a narrower range of goods. The convergence in levels was therefore progressive in incidence: the variety-deprived caught up, and they were disproportionately poor.[8] This is the extensive-margin analogue of income convergence, and the access margin of Section 2: as trade costs fell, poorer, variety-deprived economies gained access to the world's range of goods far faster than richer ones.

[7] This convergence in variety counts does not depend on the small, near-floor flows either. Dropping bilateral flows below \$50,000, the conservative screen used for the welfare gain, leaves the count-growth gradient on income pro-poor and, if anything, steeper (-0.255, correlation -0.59, against -0.191 on all flows), and the convergence-in-levels slope intact (-0.321, correlation -0.76).

[8] The income gradient is fully mediated by this channel: adding the initial variety count to the gradient regression removes it, leaving a small positive residual ($+0.047$). That conditional coefficient is uninformative about a structural income effect: initial variety is itself a mechanical determinant of the Feenstra gain and is 0.77-collinear with income, so conditioning on it partials out the convergence channel itself. It is mediation, not divergence.

4.2 The magnitude of the gains and between-country inequality

Section 4.1 measured *who* gained from variety. This section turns to how much the gains were worth in aggregate, and what they did to income inequality between countries. The variety term is a fall in the cost of living, equivalent to a gain in real consumption. The computations here use the conservative, threshold-robust variety set that drops flows below $50,000. They are partial-equilibrium cost-of-living statistics rather than general-equilibrium welfare, and they compare *across countries*, so they are largely silent on the within-country distribution studied in Fajgelbaum and Khandelwal (2016). Recall though that a partial check on within-country gains reported above suggests that we do not merely capture gains for the economic elite in poor countries.[9]

The aggregate welfare gains are sizeable. Weighting countries by population, the expansion of import variety raised real consumer welfare by an average of 15 percent over 1995 to 2023, and the gains were progressive, the poor half of countries gaining 25 percent against 12 percent for the rich half. Across the world population the progressivity compresses, with the bottom half by income gaining only somewhat more than the top, because countries with the largest populations are middle-income while the steepest gains accrue to small, very poor economies (the bottom-left panel of Figure 1 plots this global incidence curve).

Being progressive, the variety gains compressed the between-country income distribution. Aggregating them lowers the population-weighted between-country Gini of income from 0.585 to 0.567, a fall of 0.018, or 3 percent, from this single channel (the bottom-right panel of Figure 1 shows the between-country Lorenz curve shifting towards the diagonal). The effect is modest in absolute terms, a few percent of the Gini, but far from negligible if brought into context: the population-weighted between-country income Gini over 1995 to 2023 fell by about 0.13 as incomes converged, mostly through income growth in China and India, so the variety channel, at about 0.018, is equivalent to roughly a seventh of that convergence; and it comes from a source the income statistics do not cover.

These numbers understate how broad-based this channel is. Because the between-country Gini is population-weighted, China and India dominate the measured income convergence. Once these two economies are removed, between-country income convergence over 1995–2023 all but disappears, the Gini moving only from 0.509 to 0.504;[10] the variety channel is undimmed at about 0.02 points, if anything larger than for the full sample because the gains concentrate in small poor economies. For the other 136 countries, real-consumption convergence through variety was the dominant force, dwarfing a near-zero income convergence. Whatever globalisation did to inequality through the production side, on the consumer-variety margin it was a measurable force for convergence.

[9] The income gradient is much steeper for low-tier, necessity goods than for high-tier, luxury goods.

[10] This is dispersion, the population-weighted between-country Gini (σ-convergence). The broad unconditional convergence documented by Patel et al. (2021) and Kremer et al. (2022) is β-convergence in growth rates, a distinct measure (and see Johnson and Papageorgiou 2020 for a cautious survey); our point here is only that the dispersion of income did not narrow once these two economies are removed.

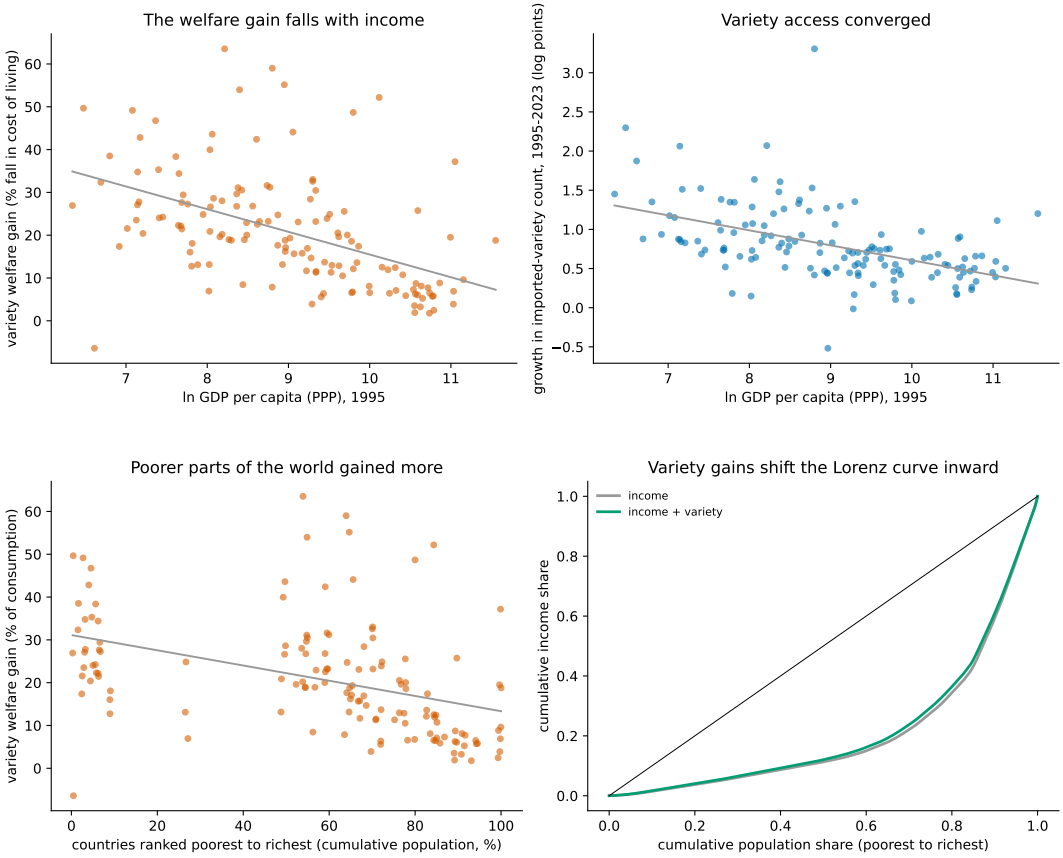

Figure 1: The variety gains were pro-poor and narrowed between-country inequality (1995–2023). *Top panel*: the Feenstra welfare gain (left) and the growth in the number of imported varieties (right) against initial income, both falling steeply with income. *Bottom panel*: a global incidence curve (left; the horizontal axis orders world population from the poorest to the richest country by income per capita) and the between-country Lorenz curve of income with and without the variety gains (right), which shift it towards the diagonal.

4.3 Variety convergence in the long run

Our results so far cover 1995 to 2023. Is variety-access convergence specific to this era of hyperglobalisation, or a longer process? The Atlas of Economic Complexity SITC-4 bilateral data, a harmonisation of UN Comtrade flows, extend the variety *counts* back to 1962, so the convergence gradient, the slope of variety-count growth on initial income, can be estimated era by era. These historical data are value-only, so this is the extensive (count) margin rather than the full welfare index.

Over the full sixty-year span the gradient is strongly pro-poor (-0.262, correlation -0.43), but Table 2 shows that this convergence is not a permanent, decade-by-decade feature of the data. Three patterns stand out. (i) The early periods yield mixed results: 1962–1971 is weakly pro-poor; the 1972–1981 decade reverses sign, with rich countries adding variety faster than poor ones in both the full and the fixed 91-country samples, a pattern consistent with the 1970s oil crises; and the 1980s are essentially flat. (ii) The strong, robust pro-poor gradient is concentrated in 1992–2001 and 2002–2011, the decades of falling tariffs, WTO

accessions and China's emergence as a major trader, where the fixed 91-country correlation reaches -0.41 and -0.59 respectively. (iii) The most recent period, 2012–2023, has barely any cross-country pattern at all, consistent with a slowdown in the cross-border decline of trade costs and a more contested trade-policy environment. The post-1995 result we uncover in the main analysis of the paper is therefore concentrated in two specific decades rather than spread evenly across a sixty-year trend.

Table 2: Variety (count) convergence by era (1962–2023)

Period	Median growth	Income Gradient (s.e.)	Corr.	N	Fixed-N (corr)
1962–2023	379%	-0.262^{***} (0.068)	-0.43	91	-0.43
1995–2023	64%	-0.115^{***} (0.020)	-0.38	139	-0.39
1962–1971	39%	-0.087^{*} (0.045)	-0.25	91	-0.25
1972–1981	10%	$+0.052$ (0.034)	$+0.15$	109	$+0.18$
1982–1991	35%	-0.022 (0.018)	-0.10	114	-0.09
1992–2001	31%	-0.071^{***} (0.017)	-0.22	139	-0.41
2002–2011	22%	-0.055^{***} (0.006)	-0.54	144	-0.59
2012–2023	7%	-0.013 (0.009)	-0.11	148	-0.10

Notes: Variety counts are from the Atlas of Economic Complexity SITC-4 bilateral data. Endpoints are three-year averages at the start and end of each time window ('era'). Median growth is the median country's proportional growth in its variety count over the time window. Gradient is the slope estimate of the growth in the number of imported (SITC-4 product) varieties regressed on initial log PPP income per capita. A negative gradient indicates pro-poor variety convergence. The correlation instead just computes the Pearson correlation between variety country growth and initial log PPP income per capita. The final column is the same Pearson correlation but with the sample restricted to the core set of countries present in the 1962–1971 window, a maximally comparable set across eras. Standard errors are heteroskedasticity-robust; *, ** and *** denote significance at the 10, 5 and 1 percent level.

5 No convergence in quality

Main results. Table 3 (Panel A) reports the quality margin for our 1995-2023 HS6 data together with those for variety welfare gain (as reported above) and variety count growth.[11]

The quality gradient is essentially flat in the full sample: the *change* in ϕ regressed on initial income gives a slope of $+0.020$ (insignificant) and a correlation of $+0.11$, barely distinguishable from zero convergence (see the top-left panel in Figure 2). The variety gradients, in contrast, are strongly progressive (correlations with initial income per capita -0.49 and -0.48), including in the variety *count* measure.[12] Poorer countries dramatically caught up in variety, but not in quality.

[11]The quality margin ϕ is usable here despite being noisier than the variety measures, because its measurement error enters the convergence regression through the *dependent* variable, where it inflates the standard error but leaves the slope unbiased.

[12]The variety count growth gradient of -0.191 and its correlation of -0.48 thus indicate stronger variety convergence in this dataset of around 5,000 products than in the SITC-4 data with merely 1,000 products (-0.115 and -0.39, respectively).

Panel B of Table 3 shows that countries which started with lower quality did not see an improvement over time. If anything, the opposite, no matter how ϕ is constructed. Across all five variants (the baseline, which screens out anomalous unit-value cells, a data-quality filter distinct from the sub-$50,000 'conservative' variety set of Section 4; a size-adjusted version that partials out a log-quantity covariate; a size-adjusted variant in combination with the data-quality filter; a three-year-pooled estimate; and a Kalman-smoothed series), the correlations stay in a narrow band from $+0.09$ to $+0.13$, several times weaker than the equivalent variety gradient (-0.49) and statistically indistinguishable from zero in every case. The contrast with variety does not depend on the particular quality index we adopt.

Yet quality did not fail to converge entirely: in a subsample of large economies, where ϕ is cleanest, it shows modest convergence, a correlation up to -0.25. But this is less than half the variety correlation for the same subsample (-0.63). The top-right panel in Figure 2 makes the contrast visible, with poor countries (orange circles) clustered at high variety gains and without a clear pattern for quality improvement. Of the 68 poorer countries, 50 are above-median in variety gains but only 30 in quality change (the grey lines mark the respective medians, counts are reported for each quadrant).

Quality is an income margin. Why did variety access converge while quality did not? Not because quality is unrelated to income, but because it is bound too tightly to income to converge any faster than income does. The flat convergence gradient and a strong dependence of quality on income are the same fact seen twice, once in changes and once in levels. Two facts, reported in Panel C of Table 3, show how tightly the quality tier is pinned to income: first, quality tracks the income *level*. Across countries ϕ rises with income at a correlation of 0.78 (bottom-left panel of Figure 2), and this is not merely a fixed national trait: in a panel regression of ϕ on log income with country and year fixed effects, the within-country slope is 0.166 (standard error 0.047, clustered by country), almost identical to the between-country slope of 0.168. A country climbs the quality tier as it grows richer. Because ϕ holds the supplier and the product fixed, this is a demand-side choice, the dearer version of the same good from the same source; whether it reflects literal quality or instead a higher willingness to pay as income rises, income is what moves it.

Second, the *change* in quality tracks income *growth*: regressing $\Delta\phi$ on income growth and initial income gives a partial correlation of 0.34 on growth and 0.22 on the initial level (bottom-right panel of Figure 2). Both patterns are stable across all five constructions of ϕ, the cross-sectional correlation ranging from 0.78 to 0.82 and the growth partial from 0.34 to 0.38 (not shown in Table 3).

Together these patterns are consistent with a simple interpretation: quality moves with income, and income barely converged over 1995 to 2023, so neither did quality. We read this as an association, not a causal estimate, and the component of income that matters is durable income. The relationship resists clean instrumentation: the exogenous shocks that

Table 3: The two margins compared and the quality margin in detail (1995–2023)

Panel A. The two margins: income gradients

Margin	Median	Gradient on income (s.e.)	Correlation
Variety welfare gain (extensive)	19.4%	-0.053^{***} (0.008)	-0.49
Variety count growth (extensive)	107%	-0.191^{***} (0.028)	-0.48
Quality tier $\Delta\phi$ (intensive)	-0.037	$+0.020$ (0.016)	$+0.11$

Panel B. Convergence of $\Delta\phi$ across constructions of ϕ

Construction of ϕ	Full sample (correlation)	Large economies (correlation, $n = 42$)
Baseline (screened)	$+0.123$ (0.079)	-0.261^{*} (0.151)
Size-adjusted	$+0.129$ (0.084)	-0.137 (0.169)
Size-adjusted, screened	$+0.106$ (0.084)	-0.160 (0.168)
Three-year pooled	$+0.091$ (0.080)	-0.273^{*} (0.146)
Kalman-smoothed	$+0.119$ (0.078)	-0.254^{*} (0.150)
For comparison:		
Variety gain convergence	-0.486^{***} (0.077)	-0.627^{***} (0.095)

Panel C. ϕ and income

Relationship	Estimate
ϕ level and income level (cross-section correlation)	$+0.776$
ϕ and income, within country (panel slope)	$+0.166^{***}$ (0.047)
ϕ and income, between country (slope)	$+0.168^{***}$ (0.012)
$\Delta\phi$ and income growth (partial correlation)	$+0.343^{***}$ (0.115)
$\Delta\phi$ and initial income (partial correlation)	$+0.220^{***}$ (0.073)
For comparison:	
Variety gain and income growth: all flows	-0.017 (0.076)
Variety gain and income growth: conservative sample	$+0.148$ (0.102)

Notes: *Panel A* reports each margin on its full HS6 sample (variety over 138 countries, quality over the 136 with a ϕ estimate); the welfare-gain row is the conservative figure in Panel A of Table 1, the count-growth row is on all flows. *Panel B* reports the correlation of $\Delta\phi$ (1995–2023) with initial income for the five constructions of ϕ; positive values indicate the absence of pro-poor convergence; "large economies" is the 42 largest, where ϕ is cleanest. *Panel C* reports the ϕ–income level and growth relationships (within-country slope clustered by country); the memo rows give the variety gain's near-zero correlation with income growth on the all-flow and conservative sets, against $+0.34$ for $\Delta\phi$. Standard errors are heteroskedasticity-robust; *, ** and *** denote significance at the 10, 5 and 1 percent level.

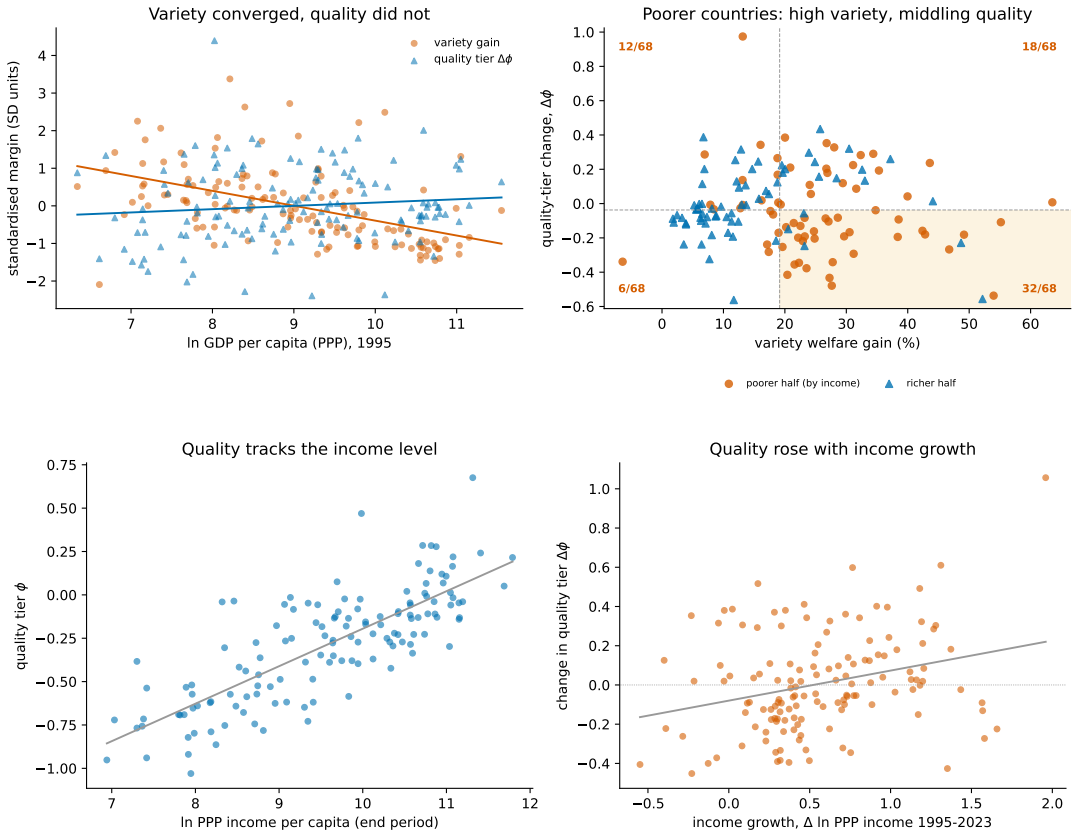

Figure 2: More choice, not more quality (1995–2023). *Top*: on a common standardised scale the variety gain falls steeply with income while the quality-tier change is flat (left); the joint distribution shows poorer countries (orange) clustered at high variety gains but split on quality (right; 50 of 68 poorer countries are above-median in variety, only 30 in quality). *Bottom*: the quality tier ϕ rises with the income level (left, correlation 0.78) and its change tracks income growth (right, partial correlation 0.34).

could instrument income, such as commodity-price or external-demand windfalls, move only its transitory component, whereas the quality tier tracks durable living standards, and durable cross-country income differences are not themselves exogenously assignable.[13] We therefore characterise the relationship by triangulation, from its appearance in cross-country levels, in within-country changes, and in long-run growth, the income-quality relationship sloping upward as non-homothetic demand predicts, the quality analogue of an Engel curve (Fieler 2011, Hallak 2006).

Supplier identity. One might be concerned that the apparent income channel masks a supplier-choice channel: imports from technologically-advanced, rich economies might carry embodied technology, in the spirit of Grossman and Helpman (1991) and Coe and Helpman (1995), and lift the importer's quality tier through *who* they buy from, rather than *how much* they earn. By construction of ϕ, this alternative explanation cannot apply here: the exporter-

[13] A direct instrumental-variables attempt, using the IMF commodity terms-of-trade index (Gruss and Kebhaj 2019) as a fixed-weight income windfall, is underpowered ($F \approx 10$) and inconclusive; full results are available on request.

by-product fixed effects $\delta_{eg,t}$ in equation (5) absorb every difference in offered quality across suppliers of the same good, so the Coe–Helpman channel that operates through supplier choice is, by design, captured by δ, not by ϕ. ϕ therefore records the tier a country selects *within* a given supplier-product cell, and it is that choice which the income relationships above explain.

6 Conclusion

We study how the variety and quality of traded goods have changed and their implications for welfare. We find that trade reached poor consumers on one margin but not the other: variety access converged broadly and progressively, a real-consumption gain worth a median of about 19 percent of the cost of living that the income statistics feeding the income convergence debate, ignoring new varieties, do not record. The quality tier did not converge: it follows income, and outside China and India income barely converged. Globalisation gave the world's poor consumers much more choice than quality.

This result adds a progressive consumer-side counterweight to the production-side account of trade's distributional record. On the variety channel alone, it lowered between-country income inequality by about 3 percent, equivalent to a sixth of the income convergence over 1995–2023 and nearly all of it outside India and China.

There are two natural extensions to our analysis. The first is a within-country treatment of the kind Fajgelbaum and Khandelwal (2016) pioneered: the necessity-versus-luxury split we apply is reassuring but only partial, and a general-equilibrium account of how the variety gains were shared *within* poor countries is an obvious next step. The second is the *supplier-side* margin in quality, which our index sets aside by design. Because ϕ holds the supplier fixed, whether importing from technologically-advanced economies raises quality through embodied knowledge is a separate question, requiring a companion index built on unit-value variation *across* suppliers rather than within.

References

Arkolakis, C., Demidova, S., Klenow, P. J., and Rodríguez-Clare, A. (2008). Endogenous variety and the gains from trade. *American Economic Review Papers & Proceedings*, 98(2), 444–450.

Arkolakis, C., Costinot, A., and Rodríguez-Clare, A. (2012). New trade models, same old gains? *American Economic Review*, 102(1), 94–130.

Arkolakis, C., Ganapati, S., and Muendler, M. A. (2021). The extensive margin of exporting products: A firm-level analysis. *American Economic Journal: Macroeconomics*, 13(4), 182–245.

Atkin, D., Faber, B., and Gonzalez-Navarro, M. (2018). Retail globalization and household welfare: evidence from Mexico. *Journal of Political Economy*, 126(1), 1–73.

Autor, D. H., Dorn, D., and Hanson, G. H. (2013). The China syndrome: local labor market effects of import competition in the United States. *American Economic Review*, 103(6), 2121–2168.

Autor, D. H., Dorn, D., and Hanson, G. H. (2016). The China shock: Learning from labor-market adjustment to large changes in trade. *Annual Review of Economics*, 8(1), 205–240.

Brambilla, I., Lederman, D., & Porto, G. (2012). Exports, export destinations, and skills. *American Economic Review*, 102(7), 3406–3438.

Broda, C. and Weinstein, D. E. (2006). Globalization and the gains from variety. *Quarterly Journal of Economics*, 121(2), 541–585.

Coe, D. T. and Helpman, E. (1995). International R&D spillovers. *European Economic Review*, 39(5), 859–887.

Costinot, A. and Rodríguez-Clare, A. (2014). Trade theory with numbers: quantifying the consequences of globalization. In *Handbook of International Economics*, vol. 4, 197–261. Elsevier.

De Loecker, J. (2013). Detecting learning by exporting. *American Economic Journal: Microeconomics*, 5(3), 1–21.

Dorn, D., and Levell, P. (2024). Trade and inequality in Europe and the US. *Oxford Open Economics*, 3(Supplement 1), i1042–i1068.

Eaton, J., Kortum, S., and Kramarz, F. (2011). An anatomy of international trade: Evidence from French firms. *Econometrica*, 79(5), 1453-1498.

Fajgelbaum, P. D. and Khandelwal, A. K. (2016). Measuring the unequal gains from trade. *Quarterly Journal of Economics*, 131(3), 1113–1180.

Feenstra, R. C. (1994). New product varieties and the measurement of international prices. *American Economic Review*, 84(1), 157–177.

Feenstra, R. C. (2018). Restoring the product variety and pro-competitive gains from trade with heterogeneous firms and bounded productivity. *Journal of International Economics*, 110, 16–27.

Fieler, A. C. (2011). Nonhomotheticity and bilateral trade: evidence and a quantitative explanation. *Econometrica*, 79(4), 1069–1101.

Gaulier, G. and Zignago, S. (2010). BACI: International trade database at the product level. The 1994–2007 version. CEPII Working Paper 2010-23.

Goldberg, P., Khandelwal, A., Pavcnik, N., and Topalova, P. (2009). Trade liberalization and new imported inputs. *American Economic Review*, 99(2), 494–500.

Grossman, G. M. and Helpman, E. (1991). *Innovation and Growth in the Global Economy*. MIT Press.

Gruss, B. and Kebhaj, S. (2019). Commodity terms of trade: a new database. *IMF Working Paper* WP/19/21.

Hallak, J. C. (2006). Product quality and the direction of trade. *Journal of International Economics*, 68(1), 238–265.

Helpman, E., Itskhoki, O., Muendler, M. A., and Redding, S. J. (2016). Trade and inequality: From theory to estimation. *The Review of Economic Studies*, 84(1), 357–405.

Hottman, C. J., Redding, S. J. and Weinstein, D. E. (2016). Quantifying the sources of firm heterogeneity. *Quarterly Journal of Economics*, 131(3), 1291–1364.

Hummels, D. and Klenow, P. J. (2005). The variety and quality of a nation's exports. *American Economic Review*, 95(3), 704–723.

Johnson, P. and Papageorgiou, C. (2020). What remains of cross-country convergence? *Journal of Economic Literature*, 58(1), 129–175.

Kremer, M., Willis, J., and You, Y. (2022). Converging to convergence. *NBER Macroeconomics Annual*, 36(1), 337–412.

Lo Forte, G. (2025). Trademarks and Gains from Variety: The Role of Multinational Enterprises. Unpublished mimeo, available at https://loforteg.github.io.

Melitz, M. J. (2003). The impact of trade on intra-industry reallocations and aggregate industry productivity. *Econometrica*, 71(6), 1695–1725.

Ossa, R. (2015). Why trade matters after all. *Journal of International Economics*, 97(2), 266–277.

Patel, D., Sandefur, J. and Subramanian, A. (2021). The new era of unconditional convergence. *Journal of Development Economics*, 152, 102687.

Pierce, J. R., and Schott, P. K. (2016). The surprisingly swift decline of US manufacturing employment. *American Economic Review*, 106(7), 1632–1662.

Rodrik, D. (2024). A primer on trade and inequality. *Oxford Open Economics*, 3(Supplement 1), i1076–i1082.