

University of Groningen

Product Variety, the Cost of Living and Welfare Across Countries

Cavallo, Alberto; Feenstra, Robert C.; Inklaar, Robert

Published in:
 American Economic Journal. Macroeconomics

DOI:
[10.1257/mac.20210137](https://doi.org/10.1257/mac.20210137)

IMPORTANT NOTE: You are advised to consult the publisher's version (publisher's PDF) if you wish to cite from it. Please check the document version below.

Document Version
 Publisher's PDF, also known as Version of record

Publication date:
 2023

[Link to publication in University of Groningen/UMCG research database](#)

Citation for published version (APA):
 Cavallo, A., Feenstra, R. C., & Inklaar, R. (2023). Product Variety, the Cost of Living and Welfare Across Countries. *American Economic Journal. Macroeconomics*, 15(4), 44-60. Article 2.
 <https://doi.org/10.1257/mac.20210137>

Copyright

Other than for strictly personal use, it is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license (like Creative Commons).

The publication may also be distributed here under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license. More information can be found on the University of Groningen website: <https://www.rug.nl/library/open-access/self-archiving-pure/taverne-amendment>.

Take-down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Downloaded from the University of Groningen/UMCG research database (Pure): <http://www.rug.nl/research/portal>. For technical reasons the number of authors shown on this cover page is limited to 10 maximum.

Product Variety, the Cost of Living, and Welfare across Countries[†]

By ALBERTO CAVALLO, ROBERT C. FEENSTRA, AND ROBERT INKLAAR*

We use the structure of the Melitz (2003) model to compute the cost of living and welfare across 47 countries, and compare these to conventional measures of prices and real consumption from the International Comparisons Project (ICP). The cost of living is inferred without directly using ICP prices of traded goods and instead relying on output prices, openness, domestic trade costs, and product variety measured by the counts of barcodes or firms. We find that welfare relative to the United States is lower than indicated by real consumption for most countries, but similar in China and Japan and similar or higher in some European countries. (JEL E21, E23, E31, O11)

Liberalizing trade is well understood to improve country welfare. Melitz (2003) is among the models that generate a very simple formula, based on openness, for the gains from trade (Arkolakis, Costinot, and Rodríguez-Clare—henceforth, ACR—2012). But while the Melitz model is well suited to compute the gains from trade *within* a country, can it be also used to compare welfare *between* countries? We will demonstrate that it can, provided that several domestic variables—product variety, domestic trade costs, and productivity—are considered. We compare the theoretical cost of living predictions from the Melitz model to *consumption price levels* from the International Comparisons Project (ICP), focusing on cross-country prices for tradable goods. We find that the theoretical cost of living relative to the United States is higher than indicated by consumption price levels for most countries, except for China, Japan, and South Korea, where the theoretical cost of living is similar to price levels, and also a select group of nations in Europe, where the cost of living is similar to or less than price levels.

*Cavallo: Harvard Business School (email: acavallo@hbs.edu); Feenstra: Department of Economics, University of California, Davis (email: rcfeenstra@ucdavis.edu); Inklaar: Faculty of Economics and Business, University of Groningen (email: r.c.inklaar@rug.nl). Simon Gilchrist was coeditor for this article. We thank seminar participants at the NBER, Georgetown University, the University of Groningen, and Yale University, along with the referees, for very helpful comments. We also thank Florencia Hnilo, Anna Ignatenko, Ninghui Li, and Mingzhi Xu for their excellent research assistance. This research is supported by the NSF (United States), SSHRC (Canada), and NWO (Netherlands) under a Digging into Data grant of the Trans-Atlantic Platform (NSF Award #1724649). Alberto Cavallo is a shareholder of PriceStats LLC, the private company that provided proprietary data used in this paper without any requirements to review the findings. Data and code for this paper are available at <https://doi.org/10.34894/7RCSFZ>.

[†]Go to <https://doi.org/10.1257/mac.20210137> to visit the article page for additional materials and author disclosure statement(s) or to comment in the online discussion forum.

Our application of the Melitz model applies to traded sectors of consumption only. To incorporate nontraded goods, we rely on those nontraded prices from the ICP (we do not collect any data on product variety or domestic trade costs in these sectors). Combining those nontraded prices with the theoretical cost of living for traded goods and using the result to deflate consumption expenditures, we obtain a theoretical measure of welfare that can be compared to real actual individual consumption, a standard statistical measure of household consumption that is not sensitive to differences in funding sources (i.e., private, nonprofit, or public) across countries. We find that welfare relative to the United States is lower than indicated by real consumption for most countries, though they are similar for certain Asian countries (including India) and similar or higher for a select group of European nations.

This paper contributes to the literature in several ways. ACR (2012) compare the equilibria in a single country facing different foreign variables, leading to different trade opportunities. The change in foreign variables can reflect differing iceberg costs of foreign trade, or trading with a foreign country of differing size (reflecting growth abroad). That approach allowed them to compute the gains from trade versus autarky, or the gains from a small change in trade. We expand on that literature by incorporating different domestic variables, meaning a change in domestic trade costs, productivity and fixed costs, and home population. That will enable us to compare equilibria between countries: in particular, differences in these variables across countries lead to differences in product variety, as we show in Section I.

Second, we demonstrate the feasibility of using online barcode data to measure product variety across countries. Barcode-level data have previously been used to compare prices across countries (Cavallo et al. 2018), and product variety to compare across cities in the United States (Handbury and Weinstein 2015) and in China (Feenstra, Xu, and Antoniadis 2020). But it is challenging to compare product variety across multiple countries because products are not identical and may be produced and sold by different firms or have differing barcode classification systems.¹ We overcome this difficulty by relying on simple counts of barcodes for large retailers in certain sectors, using online data available at the Billion Prices Project (BPP) (Cavallo and Rigobon 2016), and further estimating the share of domestic varieties using crowdsourced techniques with freelancers and a custom smartphone application. When that information could not be collected, we instead use the count of firms within each sector and country to proxy for variety counts.

Third, to compare the theoretical predictions from the Melitz model to country price levels from the ICP, we also need to control for differences in country productivity. The “next generation” of Penn World Table (PWT; see Feenstra, Inklaar, and Timmer 2015) calculates productivity using a measure of real GDP on the output side—i.e., GDP deflated with aggregate output prices that can be compared across countries. Here we use similar techniques to obtain output prices at the sectoral

¹This challenge is illustrated by Argente, Hsieh, and Lee (2022), who compare the United States and Mexico. Even though these countries share a barcode system, only 8.5 percent of the total number of Mexican barcodes and 1.5 percent of US barcodes are found in the other country, so even in this case the number of identical barcodes between the countries is quite limited. The same holds for the much broader set of countries analyzed in Beck and Jaravel (2020), where an average of 5 percent of identical barcode items are found across countries. To overcome these limitations, we rely instead on simple counts of barcodes and on the theoretical structure of the Melitz model.

level, which differ from the sectoral consumption price levels in the ICP. Therefore, our results can be used to evaluate how both the ICP and PWT datasets compare to theoretical predictions from the Melitz model.

Finally, this paper contributes to the broader literature on measures of welfare that are “beyond GDP,” to use the phrase of Jones and Klenow (2016). They propose a welfare concept that combines cross-country differences in consumption, leisure, mortality, and inequality into a single consumption-equivalent measure. Our goal here is much less ambitious: to incorporate product variety to measure the cost of living across countries while focusing on tradable goods. But there are some broad similarities in our results, particularly the finding that welfare is lower relative to real consumption in many low- and middle-income countries. The main difference is that Jones and Klenow (2016) find that nearly all Western European countries have welfare relative to the United States that is similar to or higher than indicated by conventional measures of real consumption, whereas we find that result only for a more select group of nations in Europe.

In Section I, we obtain a general expression for the theoretical cost of living (CoL) across countries in the Melitz model, which depends on domestic and exporter productivity, domestic and foreign trade costs, the terms of trade, and the extent of product variety available to consumers. Under the added assumption of a Pareto distribution for firm productivity, as in Chaney (2008), we obtain an expression that generalizes ACR (2012) to a cross-country context, where the share of expenditure on domestic goods (reflecting inverse openness) plays an important role. In Section II, we describe the data used to measure the theoretical CoL , focusing on traded goods only. In Section III, we compare that expression to the price levels for consumption (PC) across countries from the ICP. In Section IV, we add non-traded goods, and then, dividing real actual individual consumption (AIC) by CoL or PC , we can compare the theoretical measure of welfare ($U = AIC/CoL$) to real consumption ($RC = AIC/PC$) and to the results of Jones and Klenow (2016). Section V concludes, while proofs and other details are in the online Appendix.

I. Modeling Welfare between Countries

Extending the model of Melitz (2003) and Chaney (2008) to measure welfare between countries means that we must allow for differences in domestic trade costs, their productivity distributions, and populations. We model domestic trade costs as iceberg costs, meaning that in country i and sector s , $\tau_s^{ii} \geq 1$ units must be sent from the domestic firms for one unit to reach the home consumer. This is a plausible description of resources used in domestic transportation and in the wholesale and retail sectors, which we rely on to measure τ_s^{ii} .² To make the welfare comparison, we consider the ratio of equilibria in two countries that can differ in these domestic variables, as well as *foreign* variables like in ACR (2012), defined as differences in the iceberg costs of international trade and in the foreign values of population and

²Our theoretical analysis could also be extended to include differences in excise taxes across countries, which are part of our empirical analysis.

fixed costs of production and entry.³ To simplify the model, we introduce a restriction on the extent to which fixed costs can differ across countries.

The rest of the model is familiar from Melitz (2003), though extended here to multiple sectors, so our exposition will be brief. We assume a CES utility function across sectors, and a CES utility across varieties in each sector with elasticity of substitution $\sigma_s > 1$. Labor is the only factor of production, and a mass M_s^i of domestic firms pay an entry cost of F_s^i (in units of labor) to receive a productivity draw φ from the Pareto density $g_s^i(\varphi)$, with distribution function

$$(1) \quad G_s^i(\varphi) = 1 - (\varphi/A_s^i)^{-\theta} \quad \text{for } \varphi \geq A_s^i \quad \text{and } \theta > (\sigma_s - 1) > 0.$$

The lower bound for productivity A_s^i is proportional to the unconditional mean productivity, given by $\int_{A_s^i}^{\infty} \varphi g_s^i(\varphi) d\varphi = [\theta/(\theta - 1)]A_s^i$. We are allowing this mean productivity to differ across sectors and countries, but for simplicity, we are treating the Pareto parameter θ as common across sectors and countries.

With trade between country i and all countries $j = 1, \dots, C$, the CES price index for country i in sector s is defined over domestic and foreign goods as

$$(2) \quad P_s^i = \left\{ \sum_{j=1}^C M_s^{ji} \int_{\varphi_s^{ji}}^{\infty} p_s^{ji}(\varphi)^{1-\sigma_s} \frac{g_s^j(\varphi)}{[1 - G_s^j(\varphi_s^{ji})]} d\varphi \right\}^{1/(1-\sigma_s)},$$

where the consumer prices for sales from country j to i are $p_s^{ji}(\varphi)$ for variety φ , and the mass of varieties sold is $M_s^{ji} = M_s^j [1 - G(\varphi_s^{ji})]$, which adjusts the mass of entry M_s^j in country j for the probability that a firm has a profitable productivity draw $\varphi \geq \varphi_s^{ji}$ for sales to country i . We adopt a CES utility function across sectors in each country, with the elasticity η between sectors. Then the overall price index for country i is of the form

$$(3) \quad P^i = \left[\sum_{s=1}^S a_s (P_s^i)^{1-\eta} \right]^{1/(1-\eta)}, \quad \eta > 0,$$

where a_s reflect consumption weights for each sector that are common across countries.

We let λ_s^{ii} denote the share of country i expenditure on domestically produced goods. This share is obtained by taking the ratio of the term in the summation on the right of (2) for $j = i$ to the whole term in brackets, obtaining

$$(4) \quad \lambda_s^{ii} = \left\{ M_s^{ii} \int_{\varphi_s^{ii}}^{\infty} p_s^{ii}(\varphi)^{1-\sigma_s} \frac{g_s^i(\varphi)}{[1 - G_s^i(\varphi_s^{ii})]} d\varphi \right\} / P_s^{i(1-\sigma_s)}.$$

³When we consider differences in foreign variables, as in ACR (2012), the two equilibria could be for the same country facing different foreign conditions over time.

This expression can be simplified by solving for domestic prices. Country i firms face the common wage w^i , so the marginal costs of production for a firm with productivity φ is w^i/φ . Then with the usual CES markup, the domestic price is $p_s^{ii}(\varphi) = [\sigma_s/(\sigma_s - 1)](\tau_s^{ii} w^i/\varphi)$, where $\tau_s^{ii} \geq 1$ are the trade costs for domestic sales. Substituting these prices into (4) and computing the integral, we obtain

$$(5) \quad M_s^{ii} \int_{\varphi_s^{ii}}^{\infty} \left(\frac{\sigma_s}{\sigma_s - 1} \right)^{1-\sigma_s} \left(\frac{w^i \tau_s^{ii}}{\varphi} \right)^{1-\sigma_s} \frac{g_s^i(\varphi)}{[1 - G_s^i(\varphi_s^{ii})]} d\varphi \\ = M_s^{ii} \left(\frac{\theta}{\theta - \sigma_s + 1} \right) \left(\frac{\sigma_s}{\sigma_s - 1} \right)^{1-\sigma_s} \left(\frac{w^i \tau_s^{ii}}{\varphi_s^{ii}} \right)^{1-\sigma_s}.$$

Substituting (5) into (4), the share of expenditure on domestic goods is

$$(6) \quad \lambda_s^{ii} = \left(\frac{\theta}{\theta - \sigma_s + 1} \right) \left(\frac{\sigma_s}{\sigma_s - 1} \right)^{1-\sigma_s} \left(\frac{w^i \tau_s^{ii}}{\varphi_s^{ii}} \right)^{1-\sigma_s} \frac{M_s^{ii}}{P_s^{i(1-\sigma_s)}}.$$

Now compare the equilibria in sector s between countries i and j . The ratio of CES price indexes in sector s is denoted by P_s^i/P_s^j , and after we aggregate across sectors, we will obtain the theoretical *CoL* between the two countries. The sectoral price ratio is readily obtained by rearranging (6) as

$$(7) \quad \frac{P_s^i}{P_s^j} = \frac{\left[\frac{(M_s^{ii})^{\frac{1}{(1-\sigma_s)}} w^i \tau_s^{ii} / \varphi_s^{ii}}{(M_s^{jj})^{\frac{1}{(1-\sigma_s)}} w^j \tau_s^{jj} / \varphi_s^{jj}} \right] \left(\frac{\lambda_s^{ii}}{\lambda_s^{jj}} \right)^{\frac{1}{(\sigma_s-1)}}}{\left(\frac{w^i \tau_s^{ii} / \varphi_s^{ii}}{w^j \tau_s^{jj} / \varphi_s^{jj}} \right) \left(\frac{M_s^{ii} / \lambda_s^{ii}}{M_s^{jj} / \lambda_s^{jj}} \right)^{\frac{1}{(1-\sigma_s)}}}.$$

On the right-hand side of the first line of (7), the first term appearing in brackets is the ratio of the CES price index of domestic goods, where the variety term $(M_s^{ii})^{1/(1-\sigma_s)}$ is the welfare effect of any difference in the mass of domestic varieties, while $w^i \tau_s^{ii} / \varphi_s^{ii}$ is proportional to the average price of these domestic varieties (and likewise for country j). The second term is the ratio of the share of spending on domestic goods, or one minus the share of spending on imported varieties, which adjusts the price index for import varieties as in Feenstra (1994).

Either greater domestic variety in country i ($M_s^{ii} > M_s^{jj}$) or more import variety resulting in a smaller domestic share ($\lambda_s^{ii} < \lambda_s^{jj}$) reduces the relative sectoral price index in country i . By rewriting the price index as in the second line of (7), we can see that the term $M_s^{ii} / \lambda_s^{ii}$ (and likewise for country j) measures the “overall” product variety taking into account both domestic and import varieties, with an increase in overall variety lowering the price index according to the exponent $1/(1 - \sigma_s) < 0$. We describe the rest of the equilibrium conditions in online Appendix A.1.

To determine the sources of welfare differences between the two countries, we focus on the zero-cutoff-profit (ZCP) condition that determines the threshold productivity φ_s^{ii} for home domestic firms, given by

$$(8) \quad \varphi_s^{ii(\sigma_s-1)} = \frac{w^i f_s^{ii} \sigma_s}{X_s^i} \left[\frac{w^i \tau_s^{ii}}{P_s^i(\sigma_s - 1)} \right]^{\sigma_s-1},$$

where X_s^i denotes total expenditure on the differentiated good in sector s of country i , and f_s^{ii} denotes the fixed costs of production for sales to the home market. We take the ratio of (8) between two countries and substitute that into (7) to obtain

$$(9) \quad \left(\frac{M_s^{ii}/\lambda_s^{ii}}{M_s^{jj}/\lambda_s^{jj}} \right) = \left(\frac{X_s^i/w^i f_s^{ii}}{X_s^j/w^j f_s^{jj}} \right).$$

Expression (9) shows us how to solve for the “overall” product variety appearing in (7), but we still need to solve for ZCP productivity levels appearing there. As mentioned, we assume a Pareto distribution for firm productivity given by (1). The mass of firms selling domestically in country i equals $M_s^{ii} = M_s^i [1 - G(\varphi_s^{ii})] = M_s^i (\varphi_s^{ii}/A_s^i)^{-\theta}$, where M_s^i is the mass of entering firms. Then, using this in (9), we obtain

$$(10) \quad \begin{aligned} \frac{\lambda_s^{ii}}{\lambda_s^{jj}} &= \left(\frac{X_s^i/w^i f_s^{ii}}{X_s^j/w^j f_s^{jj}} \right)^{-1} \left(\frac{M_s^i}{M_s^j} \right) \left(\frac{\varphi_s^{ii}/A_s^i}{\varphi_s^{jj}/A_s^j} \right)^{-\theta} \\ &= \left(\frac{X_s^i/w^i L_s^i}{X_s^j/w^j L_s^j} \right)^{-1} \left(\frac{f_s^{ii}/F_s^i}{f_s^{jj}/F_s^j} \right) \left(\frac{\varphi_s^{ii}/A_s^i}{\varphi_s^{jj}/A_s^j} \right)^{-\theta}, \end{aligned}$$

where the final equality uses the fact that the mass of entering firms is proportional to the total labor used in sector s , which is denoted by L_s^i , so that $M_s^i \propto L_s^i/F_s^i$, as shown in online Appendix A.1.

The ratio of the fixed costs of production and entry that appears in (10) is difficult to identify from the data, so we simplify our model by assuming that it is the same across countries, which we state formally as

ASSUMPTION 1: *In each sector, the fixed costs of production for the domestic market and the fixed costs of entry are proportional across countries, $f_s^{ii}/F_s^i = f_s^{jj}/F_s^j$ for $i, j = 1, \dots, C$ and $s = 1, \dots, S$.*

With this assumption, we solve for the threshold productivity levels from (10) and use them in the ratio of the price indexes (7). Then we aggregate (7) across sectors using the CES function for the country prices in (3), which implies that the ratio of these CES prices is

$$(11) \quad \frac{P^i}{P^j} = \prod_{s=1}^S \left(\frac{P_s^i}{P_s^j} \right)^{\omega_s^{ij}},$$

with

$$\omega_s^{ij} \equiv \frac{(x_s^i - x_s^j)/(\ln x_s^i - \ln x_s^j)}{\sum_{r=1}^S (x_r^i - x_r^j)/(\ln x_r^i - \ln x_r^j)},$$

where $x_s^i \equiv X_s^i/w^i L^i$ are the sectoral expenditure shares in country i , and ω_s^{ij} are the Sato-Vartia weights.⁴ Using this aggregation, we obtain the following result:

PROPOSITION 1: *Under Assumption 1, the ratio of the CES price indexes in country i and j is*

$$(12) \quad \frac{P^i}{P^j} = \prod_{s=1}^S \left\{ \frac{w^i / \left[A_s^i (\mu_s^i)^{\frac{1}{\theta}} \right]}{w^j / \left[A_s^j (\mu_s^j)^{\frac{1}{\theta}} \right]} \right\}^{\omega_s^{ij}} \left(\frac{\lambda_s^{ii}}{\lambda_s^{jj}} \right)^{\frac{\omega_s^{ij}}{\theta}} \left(\frac{\tau_s^{ii}}{\tau_s^{jj}} \right)^{\omega_s^{ij}} \left(\frac{M_s^{ii} / \lambda_s^{ii}}{M_s^{jj} / \lambda_s^{jj}} \right)^{\frac{-\omega_s^{ij}}{(\sigma_s - 1)}},$$

where $\mu_s^i \equiv L_s^i/L^i$ denotes the share of the labor force in each sector.

To interpret this result, the first ratio on the right of (12) consists of wages relative to the sectoral productivity levels, where those productivity levels are adjusted by the labor share to each sector, μ_s^i . In practice, we will measure μ_s^i by the sectoral share of value added, reflecting the use of labor and all other factors, and therefore greater resource flows to sectors with high productivity will enhance overall productivity.

The second ratio is the share of expenditure on domestic goods, or an inverse measure of openness: if that share is lower in country i —indicating that more varieties are available from abroad—then its gains from trade are higher and the relative sectoral price in country i is lower. This “lambda ratio” is the sufficient statistic identified by ACR for the gains from trade.

The third term on the right of (12) is the ratio of domestic trade costs, so that a country with higher domestic trade costs will have correspondingly higher prices. It is surprising that the domestic trade costs do not involve an exponent reflecting the share of expenditure on domestic goods. To explain this, consider two countries where the only difference between them is that one has higher domestic trade costs, $\tau_s^{ii} > \tau_s^{jj}$. Then country i will have higher domestic prices and therefore, lower expenditure on its domestic goods, $\lambda_s^{ii} < \lambda_s^{jj}$. So, from (12), the higher domestic trade costs are partially offset by the lower domestic share, meaning that the sectoral price index does not rise in direct proportion to the higher domestic trade costs.⁵

The fourth term appearing on the right of (12) is “overall” product variety, which has a negative exponent reflecting the gains from product variety that lowers the CES relative price. Taken together, we use all the terms on the right of (12) to measure the theoretical *CoL* in country i , where we will focus on sectors producing

⁴From L'Hôpital's rule, as $x_s^i \rightarrow x_s^j$, then $\omega_s^{ij} \rightarrow x_s^j$, so the Sato-Vartia weight approaches the expenditure share.

⁵There is one parameterization, however, where the sectoral price index will rise in direct proportion to domestic trade costs, and that is where the domestic costs of transport and retail trade apply equally to domestic and imported goods. This simple case is assumed below, and in online Appendix A.2, to derive (15).

tradable goods, as described in the next section. Before turning to that empirical application, we note a final theoretical result.

Equation (9) showed how “overall” product variety in a sector is solved for in the model, based on the expenditure in that sector relative to wages times fixed costs. When we aggregate across sectors using the Sato-Vartia weights and assume trade balance so that the sum of sector expenditure equals labor income, we obtain an even simpler solution for the economy-wide “overall” product variety:⁶

$$(13) \quad \prod_{s=1}^S \left(\frac{M_s^{ii}/\lambda_s^{ii}}{M_s^{jj}/\lambda_s^{jj}} \right)^{\omega_s^{jj}} = \left(\frac{L^i}{L^j} \right) / \prod_{s=1}^S \left(\frac{f_s^{ii}}{f_s^{jj}} \right)^{\omega_s^{jj}}.$$

Thus, “overall” product variety equals the ratio of country populations divided by the sectoral average of the fixed costs of production. To interpret this result, consider for the moment a one-sector model, $S = 1$, so that the Sato-Vartia weight in (13) is equal to unity. Suppose that we compare two countries with the same population $L^i = L^j$ and the same fixed costs $f_s^{ii} = f_s^{jj}$. Substituting these conditions into (13), we immediately see that $M_s^{ii}/\lambda_s^{ii} = M_s^{jj}/\lambda_s^{jj}$. In other words, “overall” variety does not vary in the one-sector model for two countries with the same population and domestic fixed costs. This case illustrates the one-sector Melitz-Chaney model in ACR (2012), where the two equilibria being considered are in the same country (i.e., given their population and fixed costs) but are facing different foreign variables and, thus, trade opportunities: there is no welfare change across these equilibria due to changes in “overall” product variety, but only due to changes in threshold productivities φ_s^{ii} and so inverse openness λ_s^{ii} from (10). Even in a multisector model, “overall” product variety for a country with given population and fixed costs will change only due to changes in the sectoral Sato-Vartia weights in (13).

Across countries, however, we can expect to find much greater variation in “overall” product variety than within a country: (13) says that product variety will differ across countries due to their populations and their fixed costs. This result reflects the strong scale effect that operates in the Melitz-Chaney model—with larger countries having greater overall product variety—but that effect would be offset if larger countries also have higher fixed costs. An example of the link between fixed costs and country size comes from Arkolakis (2010), where firms must advertise a product to generate demand. In a simplified version of his model, advertising costs equal $f_s^{ii} = (L^i)^\alpha / \psi_s$,⁷ so that larger countries have higher fixed costs, which would limit the scale effect on variety. We do not rely this parametric form, however,⁸ and we do not restrict the fixed costs across countries except as stated in Assumption 1.⁹

⁶ While Proposition 1 relies on the assumption of a common value for θ across sectors, we show in the proof how these results can be generalized if θ is not the same across sectors, and that (13) does not rely on this assumption.

⁷ Arkolakis (2010) has a one-sector model, so that $\psi_s = \psi$. This formula simplifies his model, however, by setting another parameter $\beta = 0$ that would otherwise influence the fixed costs of advertising.

⁸ See our working paper (Cavallo, Feenstra, and Inklaar 2021), where we substitute this parametric form for fixed costs into (13) at the sectoral level and then estimate α from the resulting log-linear regression.

⁹ Note that given data on “overall” product variety on the left of (13) and on population, we can solve for the fixed costs of production on the right. In conjunction with other data on the fixed costs of entry, we can therefore begin to test Assumption 1: see online Appendix A.6.

II. The Cost of Living: From Theory to the Data

A. Sectors, Domestic Expenditures, Trade Costs, and Productivity

We shall apply (12) to measure the theoretical cost of living index across countries, although several adjustments in this equation are needed to bridge the gap between our stylized model and the data.

First, because our theory applies to traded goods, we restrict ourselves to traded sectors of consumption in this section. Specifically, we use seven sectors of consumption shown in Table 1, defined at the two-digit level of the Classification of Individual Consumption by Purpose (COICOP). For these sectors, the share of potentially traded products for household consumption varies between 100 percent (“Food, beverages and tobacco”) and 25 percent (“Other goods”).¹⁰ For example, expenditure in the “Transportation” sector includes “taxi services,” which is omitted because it is not a traded product. The domestic expenditures shares λ_s^{ii} are measured for manufactured goods in each sector s . Domestic trade costs τ_s^{ii} in sector s include the margin earned in transportation and retail trade and taxes on products, notably sales tax, value-added tax (VAT), and excise taxes. Information on the construction of both these terms is provided in online Appendix A.2, and their average values over 47 countries are reported in Table 1.

Second, our model has only labor, while in reality there are many factors of production. This feature is readily incorporated by consideration of the term w^i/A_s^i , and likewise for country j (which we choose as the United States). Let w^i denote a weighted average of factor prices used in production. The term A_s^i , the lower bound to productivity in sector s , is proportional to the mean productivity in that sector, which we measure by total factor productivity (TFP).¹¹ Using a dual approach, sectoral TFP would equal the ratio of the weighted average of factor prices to the sectoral output price. Then the ratio w^i/A_s^i would equal the output price level for the traded goods that are considered, which we denote by PY_s^{Ti} . The price level of output reflects the prices of produced goods in each country, and thus prices of exported goods are part of the price level of output (whereas import prices are included in the price level of consumption). The “next generation” of PWT (Feenstra, Inklaar, and Timmer 2015) measures the *aggregate* output price by correcting for the terms of trade in this fashion. We shall use the same approach to measure *sectoral* output prices in traded sectors (see online Appendix A.2 and Section IIIA).

Finally, letting ω_s^{Ti} equal the Sato-Vartia weight of traded goods in sector s relative to the United States, and aggregating across sectors, we can rewrite (12) as the

¹⁰Four other sectors of consumption are omitted because the products in those sectors either are all nontraded (education, hotels, and restaurants) or have very few traded products (housing and utilities, communication).

¹¹Note that the lower bound is proportional to the unconditional mean productivity, whereas TFP reflects a mean productivity that is conditional on firms surviving. Whether these are proportional could be investigated in the dynamic models of Perla and Tornetti (2014); Sampson (2016); and Perla, Tornetti, and Waugh (2021), where the lower bound and TFP both evolve over time. They assume that the lower bound of the productivity draw for new entrants is equal to the cutoff productivity for existing firms, so that under Pareto, that lower bound and the cutoff are both proportional to TFP. In a steady state, we therefore expect the lower bound and TFP to grow at the same rate, and we are assuming that the factor of proportionality between them is the same across countries.

TABLE 1—CONSUMPTION SECTORS, GOODS SHARE IN EACH SECTOR, AND VARIABLE MEANS

Sector	Code	Goods share (%)	Mean λ_s^{ii}	Mean τ_s^{ii}	Product variety
<i>Total traded consumption</i>		47			
Food, beverages and tobacco	01-02	100	0.68	1.82	Yes
Clothing and footwear	03	97	0.34	2.21	Yes
Furnishing, household equipment	05	88	0.46	2.06	Yes
Health	06	26	0.37	2.09	No
Transportation	07	57	0.49	1.87	No
Recreation and culture	09	48	0.59	1.68	Yes ^a
Other goods	12	18	0.54	2.24	Yes

Notes: Code is the COICOP code for the sector; goods share is the share of total sectoral expenditure on goods rather than services, averaged over the 47 countries.

^aBarcode data for newspapers and books are not available within the “Recreation and culture” sector.

Sources: Goods share is from World Bank (2020), λ_s^{ii} is from WIOD (2021) and OECD-WTO (2018); for a full list of sources for τ_s^{ii} , see online Appendix A.2.

theoretical cost of living for traded goods in country i relative to the United States as country j as

$$(14) \quad \frac{CoL^{Ti}}{CoL^{Tj}} \equiv \prod_{s=1}^S \left[\frac{PY_s^{Ti} (\mu_s^i)^{-\frac{1}{\theta}}}{PY_s^{Tj} (\mu_s^j)^{-\frac{1}{\theta}}} \right]^{\omega_s^{Ti}} \left(\frac{\lambda_s^{ii}}{\lambda_s^{jj}} \right)^{\frac{\omega_s^{Ti}}{\theta}} \left(\frac{\tau_s^{ii}}{\tau_s^{jj}} \right)^{\omega_s^{Ti}} \left(\frac{M_s^i / \lambda_s^{ii}}{M_s^j / \lambda_s^{jj}} \right)^{\frac{-\omega_s^{Ti}}{(\sigma_s-1)}}$$

where the weights of traded goods across sectors sum to unity, $\sum_{s=1}^S \omega_s^{Ti} = 1$. Notice that the sectoral output price in each country, PY_s^{Ti} , is weighted by the inverse of $(\mu_s^i)^{1/\theta}$, reflecting the resource use in that sector.

We stress that while (14) is measured with the data, it is still a *theoretical* cost of living implied by the Melitz model. After measuring (14), we shall compare it to the price level of consumption for each country, denoted by PC^{Ti} , based solely on the data from the 2011 round of the ICP (World Bank 2020). This price level is measured as the observed prices of tradable consumption goods in each country, converted to US dollars using the nominal exchange rate and measured relative to the US prices of the same goods.¹² By construction, then, PC^{Ti} in country i is measured relative to the United States as country j (i.e., $PC^{T,US} \equiv 1$).¹³

B. Product Variety

The last term appearing in (14) requires that we measure M_s^i , the number of domestic product varieties in each sector. Our first estimate of domestic varieties is based on the number of domestic firms active in each sector for 46 countries, taken from the Bureau van Dijk’s (2018) Orbis global dataset, which in turn is based on business registers. We eliminate duplicate names and drop firms with zero employees

¹²We use the same Sato-Vartia weights to aggregate traded-sector prices across sectors so that differences in weighting do not influence the comparison.

¹³The sectoral price of output that appears in (14), PY_s^{Ti} , uses some of the same disaggregated prices collected across countries used to construct the sectoral price level of consumption, PC_s^{Ti} . But as we have already noted, the price of output incorporates export prices, while the price of consumption goods incorporated import prices, so these two data series differ by the terms of trade as discussed in Section IIIA.

to eliminate shell companies. As a verification exercise, we also collected data on the number of firms from national enterprise statistics, primarily from the OECD Structural Business Statistics and Eurostat Enterprise Statistics, supplemented by national reports.¹⁴ For most countries, the correspondence between the two sources is close; the correlation of the log number of firms between both sources is 0.75, rising to 0.90 when excluding India and Indonesia. The reason India and Indonesia are outliers is their large informal sector. The firm count from the enterprise statistics includes informal firms, while Orbis only counts formal firms, and the presence of a large informal sector can reduce opportunities for formal firms.

The scale effect of country size on product varieties is illustrated in the left panel of Figure 1, where we find a strong correlation of 0.82 between the number of firms in Orbis and the population in each country (in logs). Once again, India (IND) and Indonesia (IDN) are outliers due to the large informal sectors not accounted for in the Orbis data.¹⁵

The number of firms is a very crude measure of the number of products, however, because of multiproduct firms, for example. Large consumer firms sell many products, while in the firm count data, each firm is only counted once. At the other extreme, certain low-income countries like India have more fragmented markets, with some firms serving only a single city or neighborhood. We might think of these firms as producing less than a single (national) product. In addition, the count of firms applies to those producing both final goods and intermediate inputs in each sector, whereas our theory applies to the product variety of final goods only.

To obtain a more accurate count of product variety, we rely on barcode counts for goods sold within each of the sectors shown in Table 1 except for “Health” and for “Transportation.”¹⁶ These barcode counts are obtained from microdata available at the BPP (Cavallo and Rigobon 2016; Cavallo, Feenstra, and Inklaar 2022) for all products sold by some of the largest multichannel retailers in 24 countries: Australia, Brazil, Canada, Chile, China, Colombia, France, Germany, Greece, India, Ireland, Italy, Japan, Mexico, the Netherlands, New Zealand, Poland, Russia, South Africa, South Korea, Spain, Turkey, the United Kingdom, and the United States. The data were collected on a daily basis from the websites of the largest retailers in each country by PriceStats, a private company related to the BPP. PriceStats uses scraping software to automate the data collection and cleans and categorizes all individual products at a four-digit COICOP level. The collected product barcodes are usually stock-keeping units that are unique to each product and retailer.

To compute the barcode counts at the sector level, we first take the daily barcode count for each retailer in the BPP sample during 2018, calculated at a four-digit COICOP level of aggregation (e.g., “coffee, tea, and cocoa”). Next, we take the mode to control for days with no data or other outliers, giving us a single number that corresponds to the most common daily barcode count for each retailer and

¹⁴ See online Appendix A.2 for a full list of sources.

¹⁵ The data for firm counts and barcode counts, discussed next, are shown in online Appendix Tables A1–A4.

¹⁶ For “Health” and for “Transportation,” we do not have barcodes, so we continue to use the firm counts. In addition, barcodes for the “Recreation and culture” sector are not available for all products in that sector (see the notes to Table 1), so λ_s^{ii} is adjusted to match this coverage based on data from WIOD (2021) and OECD-WTO (2018).

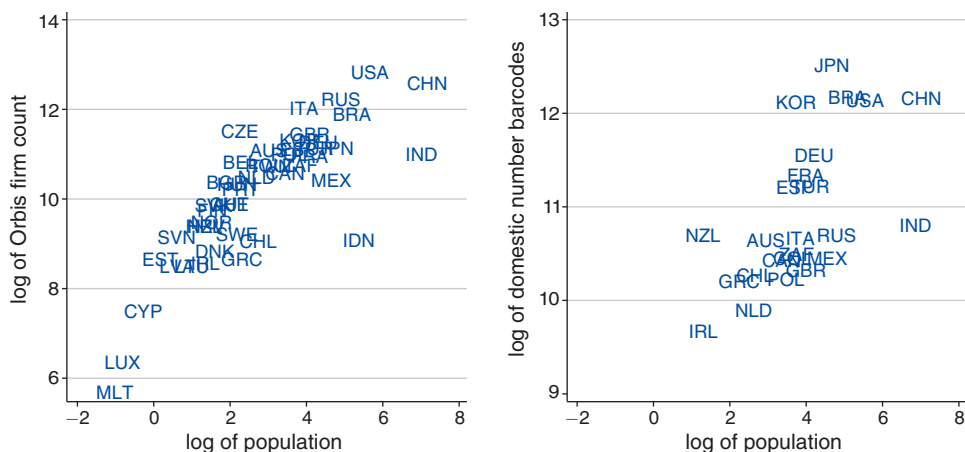


FIGURE 1. COUNT OF FIRMS AND DOMESTIC BARCODES VERSUS POPULATION (LOG SCALE)

Notes: The count of firms is obtained for seven sectors (described in Table 1) and summed to obtain the results in the left panel. The count of domestic barcodes is obtained for five sectors (omitting “Health” and “Transportation”) and summed to obtain the results in the right panel.

Sources: Firms counts from Bureau van Dijk (2018); see online Appendix A.3. Barcode counts developed in this paper; see Section IIB and online Appendix A.3. Population is from PWT 10.0 (Feenstra, Inklaar, and Timmer 2015, 2021).

category combination. To avoid double counting identical products sold by different retailers, we keep the largest barcode count for each four-digit category available across retailers and add up all the barcodes to get a count for the sectoral level (e.g., “Food, beverages and tobacco”). The number of barcodes in each sector is denoted by N_s^i . We note that all retailers included in the BPP sample have large market shares of either total or sectoral retail sales in their own country. In some countries, there are a handful of retailers that dominate the retail market, selling the largest number of varieties in most categories. In other countries, retail sales can be spread among specialized retailers that dominate a single sector. By keeping the largest barcode count in each four-digit category, we are therefore obtaining more comparable samples across countries, representing the maximum number of varieties that a consumer can find at that level by visiting any given retailer.

When collecting the count of barcodes N_s^i in each sector, we are including both domestically produced and imported goods. To measure M_s^i , we need to estimate the number of *domestic* varieties. So for two sectors—“Food, beverages and tobacco” and “Recreation and culture”—we further collected the country-of-origin information for a random sample of the total number of barcodes.¹⁷ Specifically, we hired 56 freelancers in 19 countries to manually check 500 randomly sampled barcode items per sector in each country. Using a custom mobile phone application, each freelancer visited one of the retailers in the BPP sample, scanned the barcode of each product, took a photo of the country-of-origin label, and determined whether the product was domestic or imported. When no country of origin is listed, then the product is treated

¹⁷ Within “Recreation and culture,” we collected country-of-origin information for barcodes in electronics and certain other consumer products such as bicycles.

as domestically made. The resulting dataset contains over 16,000 barcodes collected from 100 retailers, with more details provided in online Appendix A.3. For these two sectors, we therefore can calculate the barcode domestic ratio—i.e., the ratio of domestically produced barcodes to the total number of sampled barcodes, which is denoted by B_s^i .¹⁸ The number of domestically produced barcodes is therefore $M_s^i = N_s^i B_s^i$. In addition, for these sectors we also have the expenditure domestic ratio, which we have denoted by λ_s^{ii} . The overall measure of product variety is therefore

$$\left(\frac{M_s^i}{\lambda_s^{ii}}\right) = \left(\frac{N_s^i B_s^i}{\lambda_s^{ii}}\right).$$

Outside of “Food, beverages and tobacco” and “Recreation and culture” (and for some countries within those sectors),¹⁹ we do not have information on the share of domestically produced barcodes. In these cases, we make the simple assumption that $B_s^i \approx \lambda_s^{ii}$, so that the number of domestic barcodes is computed as $M_s^i = N_s^i \lambda_s^{ii}$. That is, we are assuming that the domestic share of barcodes is equal to the domestic share of expenditure. This simple assumption is plausible given that for the countries where we collect the country-of-origin information, the median value of the domestic share B_s^i for “Food, beverages and tobacco” is 0.78, compared to a median value of 0.73 for λ_s^{ii} , while in “Recreation and culture,” those medians are 0.18 and 0.14, respectively.

In the right panel of Figure 1, we show the number of domestic barcodes M_s^i for 24 countries, summed over the 5 sectors for which we have barcode counts (i.e., excluding “Health” and “Transportation”).^{20,21} There is a correlation of 0.61 with the log of population, again with some outliers such as India, where the large informal sector likely reduces the variety that can be captured with online data and crowdsourcing methods in formal retailers.²² For the countries with both firm and domestic barcode counts, we find a correlation of 0.63 between the number of firms and the number of domestic varieties.

¹⁸To test the validity of our estimates for the barcode domestic ratios, we also computed an alternative metric using country-of-origin information collected online for individual products in a subset of nine countries. This information was scraped from the website of a single retailer in each country. The resulting dataset has more products but fewer retailers. The correlation between the online and offline barcode domestic ratios is 0.76. We also found similar ratios for food in the United States using Nielsen’s scanner data. More details are provided in online Appendix A.4.

¹⁹No data for either sector could be collected for Chile, New Zealand, South Africa, or South Korea. No data for “Recreation and culture” could be collected for India and Italy.

²⁰Barcode-count data cover 26 countries including Argentina and Uruguay, but we do not have data on the share of domestic expenditure for these two countries. We also lack data on the number of firms from Orbis for Colombia, which means our analysis based on this variety measures covers 46 countries. See online Appendix Table A1 for the firm counts in 46 countries and 7 sectors, and online Appendix Table A2 for the barcode counts in the 24 countries that we cover in this paper. Focusing on the common five sectors (i.e., excluding “Health” and “Transportation”) and 23 countries (excluding Colombia), online Appendix Table A4 shows that the number of firms exceeds the number of domestic barcodes in 10 countries, and is less than the number of barcodes in 13 countries. Having more firms than barcodes in these sectors can occur because some of the firms might exclusively produce intermediate inputs, while having more barcodes than firms can occur because of multiproduct firms.

²¹Barcode data and domestic shares are available as part of the Data and Code for this paper, available at <https://doi.org/10.34894/7RCSFZ>.

²²India’s online retail sector was relatively undeveloped when these barcodes were counted in 2018. The country ranked last in the UN “e-commerce index” among those included in our barcode sample (UNCTAD 2017) and the World Bank estimated that only about 1.6 percent of sales took place online that year (Kathuria et al. 2019).

C. Parameter Values

Also needed in (14) are the elasticity of substitution σ_s and the Pareto parameter θ , which we treat as the same across sectors. For the Pareto parameter, we obtained an estimate for the Melitz-Chaney model by relying on the simulated method of moments from Simonovska and Waugh (2014), who also use cross-country data on the prices of goods collected by the ICP.²³ Using ICP data from 2011 (World Bank 2020), we obtain the pooled estimate of $\theta = 5.1$ across all sectors.

For the elasticities of substitution, using firm counts to measure product variety suggests using the elasticity of substitution across firms from Hottman, Redding, and Weinstein (2016). While our theory allows for differing values of σ_s across sectors, we consider their median elasticity across firms of 3.9.²⁴ Comparing this estimate to $\theta = 5.1$ gives a ratio for $\theta/(\sigma - 1)$ of 1.75. Note that Eaton, Kortum, and Kramarz (2011, 1472) also find an initial estimate of 1.75 for this ratio from French data on exporting firms. So $(\sigma, \theta) = (3.9, 5.1)$ is our initial, low set of estimates for the parameters.

A set of higher-parameter estimates comes from considering the elasticity of substitution across barcode items (within firms) from Hottman, Redding, and Weinstein (2016), which has a median value of 6.9. That estimate is quite close to the median elasticity from barcode items in Redding and Weinstein (2020), which is $\sigma = 6.5$.²⁵ That value can be compared to the initial estimate of $\theta = 8.3$ from Eaton and Kortum (2002, 1754). The ratio of these parameters gives $\theta/(\sigma - 1) = 1.5$, which is slightly lower than found just above but still an acceptable spread between the parameters. So, this alternative approach gives us high estimates of $(\sigma, \theta) = (6.5, 8.3)$. We have made all our calculations using both the low and high sets of estimates and find that the results are quite similar. The range of the cost of living estimates is about 1.5–2 times greater when the low-parameter estimates are used, but the relative position of countries is much the same. We present in the text the results obtained with the high estimates $(\sigma, \theta) = (6.5, 8.3)$ and report in online Appendix A.4 the results obtained using the low estimates $(\sigma, \theta) = (3.9, 5.1)$.

III. Empirical Results on the Cost of Living

In this section we show our empirical results for the cost of living for traded goods in equation (14). We first focus on the terms related to openness, domestic trade costs, and the output price, and later add the effects of the product variety.

²³We are grateful to Mike Waugh for providing us with the programs required to run these estimates; see <https://github.com/mwaugh0328>.

²⁴We found that using estimates for σ_s that vary across sectors will exaggerate the effects of product variety differences across countries in those sectors with low values of the elasticity.

²⁵Both Hottman, Redding, and Weinstein (2016) and Redding and Weinstein (2020) estimate elasticities of substitution across barcode varieties using the Nielsen Homescan data. These barcodes are for grocery store items, and many of their barcodes are within our “Food, beverages, and tobacco” sector.

A. Openness, Domestic Trade Costs, and the Terms of Trade

We refer to the lambda ratio $(\lambda_s^{ii}/\lambda_s^{jj})^{1/\theta}$ that appears in (14) as “inverse openness.” Both this variable and the ratio of domestic trade costs $(\tau_s^{ii}/\tau_s^{jj})$ are constructed at the sectoral level, with $j = USA$. For convenience in graphing these, however, we take the weighted average across sector within each country, to obtain $\prod_{s=1}^S (\lambda_s^{ii}/\lambda_s^{jj})^{\omega_s^{Ti}/\theta}$ and $\prod_{s=1}^S (\tau_s^{ii}/\tau_s^{jj})^{\omega_s^{Ti}}$. In Figure 2 we plot these variables against the price level of tradable consumption goods, PC^{Ti} , in natural logs with the United States at point (0,0) in both panels. As noted earlier, we construct PC^{Ti} by aggregating the prices of tradable consumption goods within each country from the 2011 round of the ICP using Sato-Vartia weights. So, in contrast to the theoretical cost of living in (14), which we are measuring using the Melitz model, this consumption price level simply reflects price data collected across countries and expressed relative to the United States.

The first panel in Figure 2 show that inverse openness is negatively correlated with the PC^{Ti} , implying that openness increases with the price of consumption captured by ICP. This result seems counterintuitive, but note that many of the countries that are more open than the United States are also countries with higher domestic trade costs, as shown by the positive correlation between domestic trade costs and PC^{Ti} in the second panel. Denmark (DNK) and the Netherlands (NLD), in particular, are two countries with the highest openness in the first panel (lowest inverse openness) and also the greatest domestic trade costs in the second panel, contributing to high consumption prices.

Next, we calculate all the factors in the cost of living in (14)—including the output price, domestic trade costs, and openness—*except* for variety. In Figure 3, the first panel plots the log of $\prod_{s=1}^S \left\{ \left[PY_s^{Ti} (\mu_s^i)^{-1/\theta} \right] / \left[PY_s^{Tj} (\mu_s^j)^{-1/\theta} \right] \right\}^{\omega_s^{Ti}} (\lambda_s^{ii}/\lambda_s^{jj})^{\omega_s^{Ti}/\theta} \times (\tau_s^{ii}/\tau_s^{jj})^{\omega_s^{Ti}}$ against the consumption price of traded goods, PC^{Ti} . The cost of living (without variety) and the consumption prices are tightly clustered around the 45-degree line in the first panel of Figure 3. From the second panel, we see that slightly more countries have a cost of living (without variety) relative to the United States that is lower versus higher than indicated by their consumption price level. Those with a cost of living below their relative consumption price include Russia, some countries in Eastern Europe (Croatia, HRV; Czechia, CZE; Estonia, EST; Hungary, HUN; Poland, POL; Slovakia, SVK; and Slovenia, SVN) and many in Western Europe (Austria, AUT; Belgium, BEL; Denmark, DNK; France, FRA; Germany, DEU; Ireland, IRL; Luxembourg, LUX; the Netherlands, NLD; Spain, ESP; Sweden, SWE; Switzerland, CHE; and the United Kingdom, GBR), as well as Canada (CAN) and Mexico (MEX). Denmark, Luxembourg, and the Netherlands are extreme cases where the cost of living is more than 20 percent below their relative consumption prices.

To understand these differences, we need to study the impact of the output price levels PY_s^{Ti} . These output prices differ from the consumption prices PC_s^{Ti} because consumption prices include imports, whereas output prices include exports. These two variables therefore differ by the terms of trade. The terms of

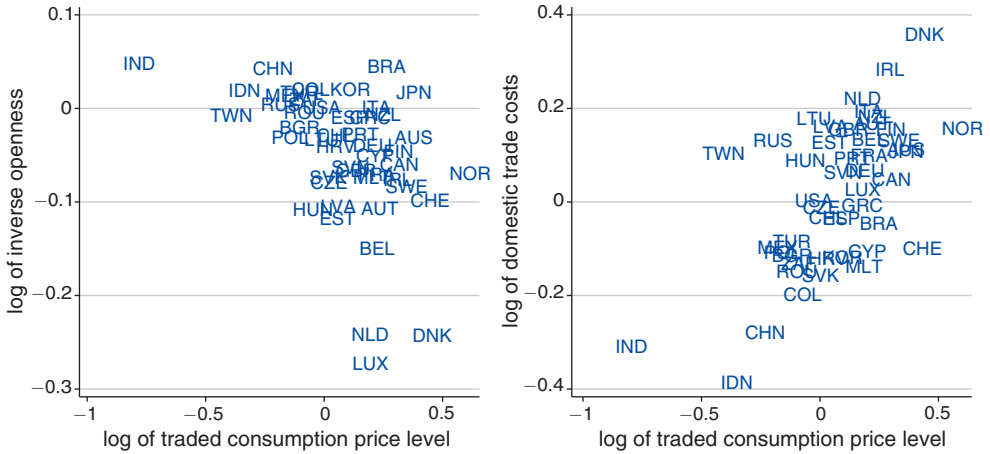


FIGURE 2. INVERSE OPENNESS AND DOMESTIC TRADE COSTS VERSUS TRADED CONSUMPTION PRICE LEVEL (LOG SCALE, USA = 0)

Notes: This figure plots inverse openness (left panel) and domestic trade costs (right panel) against the price level of tradable consumption goods. All variables are expressed in natural logs with the United States at point (0,0) in both panels.

Sources: See notes to Table 1 for sources on (inverse) openness and domestic traded costs. Traded consumption price levels are computed based on World Bank (2020).

trade are constructed from the quality-adjusted export and import prices estimated by Feenstra and Romalis (2014). As already mentioned, these prices are used in PWT to construct an aggregate output price, and here we follow much the same procedure to construct sectoral output prices. Specifically, to obtain the output prices, we start with the price level of consumption for traded goods and net out the domestic trade costs (which we assume are identical for domestically produced and imported goods); then we add export prices; and finally, we net out tariff-inclusive import prices. This calculation (see online Appendix A.2) gives

$$(15) \quad \frac{PY_s^{Ti}}{PY_s^{Tj}} = \left(\frac{PC_s^{Ti} / \tau_s^{ii}}{PC_s^{Tj} / \tau_s^{jj}} \right)^{\frac{1-\omega_s^{Xi}}{1-\omega_s^{Mi}}} \left(\frac{\tilde{P}_s^{Xi}}{\tilde{P}_s^{Xj}} \right)^{\omega_s^{Xi}} \left(\frac{\tilde{P}_s^{Mi}}{\tilde{P}_s^{Mj}} \right)^{-\frac{\omega_s^{Mi}(1-\omega_s^{Xi})}{1-\omega_s^{Mi}}},$$

where \tilde{P}_s^{Xi} and \tilde{P}_s^{Mi} are quality-adjusted prices for exports and imports,²⁶ while ω_s^{Xi} and ω_s^{Mi} are the associated Sato-Vartia weights. Substituting (15) into (14), we

²⁶These prices are obtained from Feenstra and Romalis (2014), where the quality-adjusted import price is measured net of tariffs since it is used to deflate duty-free imports in GDP. But here we measure it inclusive of tariffs.

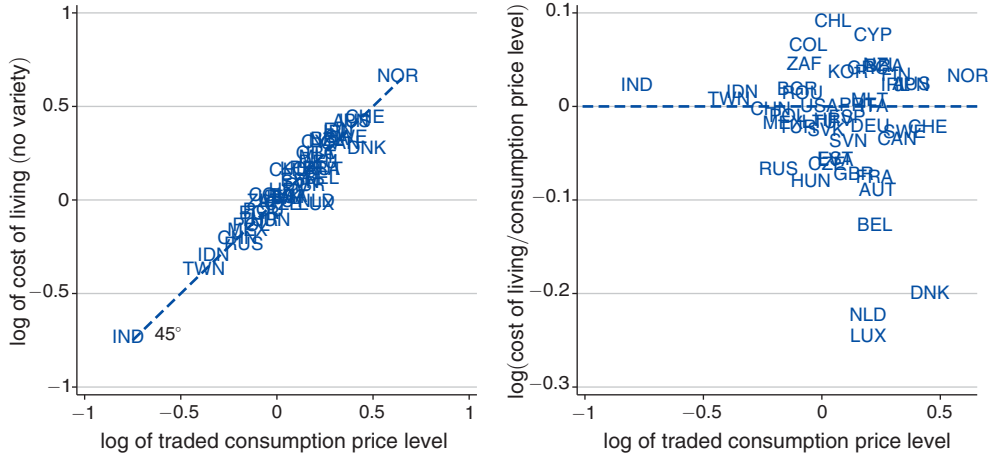


FIGURE 3. COST OF LIVING (WITHOUT VARIETY) DUE TO PRODUCTIVITY, OPENNESS, AND THE DOMESTIC TRADE COSTS VERSUS TRADED CONSUMPTION PRICE LEVEL (LOG SCALE, USA = 0)

Notes: The left-hand panel plots $\ln CoL^{Ti}$ versus $\ln PC^{Ti}$ for the 47 countries in our analysis, with $\ln CoL^{Ti}$ as defined in equation (14) (excluding the variety effects) and $\ln PC^{Ti}$ computed as the price level of traded consumption, normalized to USA = 1. The right-hand panel plots $\ln(CoL^{Ti}/PC^{Ti})$ versus $\ln PC^{Ti}$.

obtain²⁷

$$(16) \quad \frac{CoL^{Ti}}{CoL^{Tj}} \equiv \prod_{s=1}^S \left(\frac{PC_s^{Ti}}{PC_s^{Tj}} \right)^{\frac{\omega_s^{Ti}(1-\omega_s^{Xi})}{(1-\omega_s^{Mi})}} \left(\frac{\tau_s^{ii}}{\tau_s^{jj}} \right)^{\frac{\omega_s^{Ti}(\omega_s^{Xi}-\omega_s^{Mi})}{(1-\omega_s^{Mi})}} \left(\frac{\tilde{P}_s^{Xi}}{\tilde{P}_s^{Xj}} \right)^{\omega_s^{Ti}\omega_s^{Xi}} \left(\frac{\tilde{P}_s^{Mi}}{\tilde{P}_s^{Mj}} \right)^{\frac{\omega_s^{Ti}\omega_s^{Mi}(1-\omega_s^{Xi})}{(1-\omega_s^{Mi})}}$$

$$\times \prod_{s=1}^S \left(\frac{\mu_s^i}{\mu_s^j} \right)^{-\frac{\omega_s^{Ti}}{\theta}} \left(\frac{\lambda_s^{ii}}{\lambda_s^{jj}} \right)^{\frac{\omega_s^{Ti}}{\theta}} \left(\frac{M_s^i/\lambda_s^{ii}}{M_s^j/\lambda_s^{jj}} \right)^{-\frac{\omega_s^{Ti}}{\sigma-1}}.$$

To interpret the first line of (16), consider the simplified case where trade is balanced sector by sector, with $\omega_s^{Xi} = \omega_s^{Mi}$. In that case the first line starts with the weighted consumption price level. The next term, which is domestic trade costs, disappears when $\omega_s^{Xi} = \omega_s^{Mi}$ because it equally impacts the cost of living (on the left) and the consumption price level (on the right). The remaining terms on the first line are interpreted as the terms of trade—i.e., the price of exports relative to imports.

In the Melitz model, the beneficial impact of trade is measured by openness, which lowers the cost of living by appearing inversely on the second line of (16), but in ICP or PWT data, the beneficial impact of trade is measured by the terms of trade, which lower the consumption price level as compared to the output price level. In Figure 4, we plot openness (i.e., $-\ln \prod_{s=1}^S (\lambda_s^{ii}/\lambda_s^{jj}) \omega_s^{Ti}/\theta$) against the terms of trade (i.e., the log of the term on the first line of (16) except for the consumption price

²⁷ In online Appendix Tables A5 and A6, we show the log values of the terms in (16) to provide a decomposition of the cost of living relative to the consumption price level.

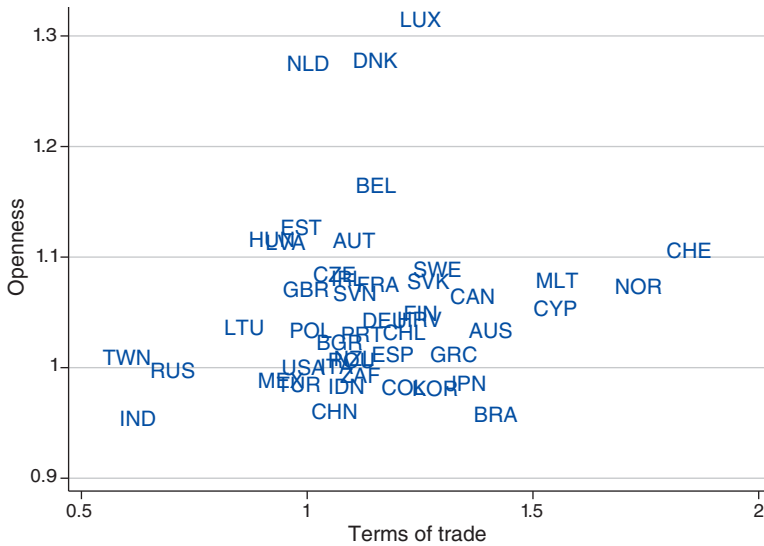


FIGURE 4. OPENNESS VERSUS THE TERMS OF TRADE BY COUNTRY (USA = 1)

Note: This figure plots the openness of countries against their terms of trade, as defined in the main text.

level and domestic trade costs). We see that there is a positive correlation between the two, though with some outliers.

Consider Denmark, Luxembourg, and the Netherlands in Figure 4, which are very open but have terms of trade that are not much different than those for the United States—i.e., close to unity. Their openness contributes to a low cost of living relative to the United States, while having terms of trade close to unity *does not* contribute to low consumption prices. As a result, these three countries have the lowest costs of living (without variety) as compared to consumption prices in Figure 3. Then consider Switzerland (CHE), which is somewhat more open than the United States but has the highest terms of trade in Figure 4, which contributes to low consumption prices. As a result, its cost of living as compared to its consumption price is higher (i.e., close to USA = 0) in Figure 3. We conclude that openness versus the terms of trade contributes meaningful variation to the cost of living (without variety) relative to the consumption price level.

B. Product Variety

The main novelty of our approach is to incorporate product variety. As previously described, we will be using two measures of product variety: the count of barcodes and the count of firm. We should recognize, however, that the count of firms is potentially a proxy for “true” product variety. Specifically, suppose that the firm-count measure of product variety \tilde{M}_s^i , by sector s , is related to “true” variety M_s^i measured using the barcode count according to

$$(17) \quad \ln \tilde{M}_s^i = \beta_s + \beta \ln M_s^i + \varepsilon_s^i.$$

Differencing with respect to the United States as country j , and taking the weighted average across sectors to reduce the errors, we obtain

$$(18) \quad \sum_{s=1}^S \omega_s^{Ti} \ln(\tilde{M}_s^i / \tilde{M}_s^j) = \beta \sum_{s=1}^S \omega_s^{Ti} \ln(M_s^i / M_s^j) + \varepsilon^{ij},$$

with the error $\varepsilon^{ij} \equiv \sum_{s=1}^S \omega_s^{Ti} \ln(\varepsilon_s^i / \varepsilon_s^j)$.

The ordinary least squares estimate of (18) gives $\hat{\beta} = 2.03$ (SE = 0.24). We conclude that taking approximately the square root of the firm count gives an estimate of variety that is reasonably close to that obtained from barcode count (for those countries where we have both sources of data).²⁸ Accordingly, we move $\beta \approx 2$ to the left of (18), and the variety effect in (16) becomes $\sum_{s=1}^S \frac{\omega_s^{Ti}}{\sigma-1} \ln\left[\left(\sqrt{\tilde{M}_s^i} / \lambda_s^{ii}\right) / \sqrt{\tilde{M}_s^j} / \lambda_s^{jj}\right]$ when measured with firm count. This is compared to the “overall” variety effect using the barcode count in Figure 5. It is clear that these two measures are highly correlated with very similar scales, and that most—but not all—countries have less product variety than the United States.²⁹ From now on, we use approximately the square root of the firm count when measuring variety with those data.³⁰

We combine the barcode and firm count datasets by using the count of barcodes to measure product variety for the 24 countries for which those data are available and by using approximately the square root (or $1/\hat{\beta} = 0.49$) of the firm count for the other 23 countries for which barcodes are not available, for a total sample of 47 countries. Figure 6 shows the estimates of the cost of living—including product variety—versus the consumption price level. The 23 countries where the firm counts are used have a 95 percent confidence band around the cost of living, because of the error in estimating $\hat{\beta}$, while the countries where barcode data are used do not have a confidence interval.³¹

The variety effect increases the cost of living in all countries relative to the United States, which has nearly the greatest variety. As a result, most countries have a greater cost of living relative to the United States than indicated by their relative

²⁸In online Appendix A.3, we discuss what features of the firm and barcode counts lead to this “square root rule.” We show that the firm count in a very small country like Ireland is roughly 100 times lower than in the United States, whereas the count of domestic barcodes is only 10 times lower. That pattern is repeated for other small countries like New Zealand, leading to the square root relation. Evidently, the surviving firms in these small and very open countries have more product varieties per firm, on average, than they do in a large and less open country like the United States (as suggested by the theoretical results in Feenstra and Ma 2009).

²⁹Our “overall” measure of product variety in Figure 5 is measured by the last term in (16), including its negative exponent, and is weighted across sectors. This figure shows that only the Netherlands has greater “overall” product variety than the United States when measured using firm counts, by appearing below the USA = 1 point, while Germany, Japan, the Netherlands, and South Korea have slightly greater overall variety than the United States when using barcode counts.

³⁰To be precise, we use $1/\hat{\beta} = 0.49$ to transform the firm count, and we also take into account the standard error of this estimate to compute the 95 percent confidence intervals in Figures 6 and 7 for countries where we use firm counts.

³¹In online Appendix Figures A2 and A3, respectively, we show the cost of living results separately for the 46 countries using firm counts (which excludes Colombia, for which firm counts are not available) to measure product variety and for the 24 countries using barcode counts. As suggested by Figure 5, the overall variety effects using barcodes or firm counts are similar enough that Figure A2, using entirely firm counts, has much the same pattern as Figure 6, using mixed barcodes and firm counts. The main difference is for Russia, which has more firms relative to the United States than barcodes, and so its theoretical cost of living is lower when using the firm counts.

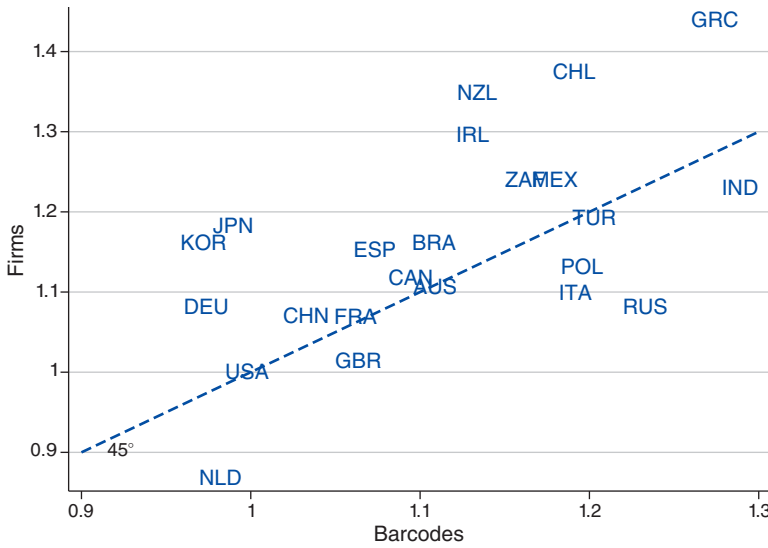


FIGURE 5. VARIETY EFFECTS BY COUNTRY—FIRM COUNT VERSUS BARCODE COUNT (USA = 1)

Note: This figure plots the aggregate variety effect relative to the United States.

consumption prices (second panel). Only Belgium, Denmark, Germany, Luxembourg, and the Netherlands have a cost of living relative to the United States that is lower than their traded consumption price level. A group of other countries have relative costs of living that are not substantially different from their relative consumption prices. This group includes several countries in Europe: Austria, Czechia, France, Hungary, and the United Kingdom (GBR), along with China, Japan, and South Korea.

C. Variance Decomposition

To further examine the relationship between CoL^{Ti} and PC^{Ti} , and to understand how the different factors contribute to their difference, we perform a decomposition of variance like that in Eaton, Kortum, and Kramarz (2004). We take the difference between the “true” cost of living in (16) and the price of consumption,

$$(19) \quad \Delta \ln \left(\frac{CoL^{Ti}}{CoL^{Tj}} \right) \equiv \ln \left(\frac{CoL^{Ti}}{CoL^{Tj}} \right) - \ln \left(\frac{PC^{Ti}}{PC^{Tj}} \right),$$

corresponding to the second panel in Figure 6. The log of all the terms appearing on the first line of (16) are denoted by $\ln(Z_i^i/Z_1^i)$, which we refer to as “trade costs plus the terms of trade,” since they include tariffs (in the import prices), domestic trade costs (when $\omega_s^{Xi} \neq \omega_s^{Mi}$), and the terms of trade. The other terms appearing on the second line of (16) are denoted by $\ln(Z_k^i/Z_k^i)$, $k = 2, 3, 4$, which refer to the sectoral shares, inverse openness, and overall variety.

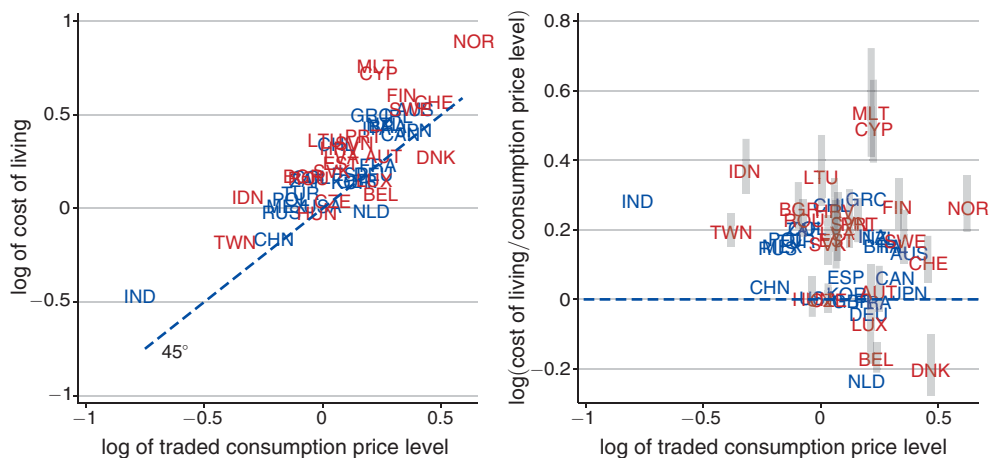


FIGURE 6. COST OF LIVING (WITH PRODUCT VARIETY) VERSUS THE TRADED CONSUMPTION PRICE LEVEL
(LOG SCALE, USA = 0)

Notes: The left-hand figure plots $\ln CoL^{Ti}$ versus $\ln PC_c^{Ti}$ for the 24 countries with barcode counts (shown in blue) and the 23 countries with only firm counts (shown in blue, with confidence intervals). The variable $\ln CoL^{Ti}$ is defined in equation (14) and $\ln PC_c^{Ti}$ computed as the price level of traded consumption, with PC^{Ti} and CoL^{Ti} normalized to USA = 1. The right-hand figure plots $\ln(CoL^{Ti}/PC^{Ti})$ versus $\ln PC^{Ti}$.

We define $\Delta \ln(Z_k^i/Z_k^j) \equiv \ln(Z_k^i/Z_k^j) - \ln(PC_k^i/PC_k^j)$ as the difference with the consumption price level, and we run the regressions:

$$(20) \quad \Delta \ln\left(\frac{Z_k^i}{Z_k^j}\right) = \gamma_{0k} + \gamma_k \Delta \ln\left(\frac{CoL^{Ti}}{CoL^{Tj}}\right), \quad k = 1, \dots, 4.$$

These regressions aim to account for the cross-country variation in the *difference* between the relative cost of living and the consumption price level. Table 2 presents the results. By construction, the regression coefficients shown in Table 2 sum to unity, so we can interpret them as the portion of the variation in the cost of living difference relative to the consumption prices.

In column (1), we ignore the overall product variety term, so we construct the cost of living as in (16) but without the final variety term. In that case, three variables explain the cost of living (without variety): trade costs plus the terms of trade, the sectoral shares, and inverse openness. Using each of these as a dependent variable in the regression (20), we see that the trade costs plus the terms of trade explain about 11 percent of the variance in the cost of living relative to the consumption price level, the sectoral shares explain 3 percent, and inverse openness explains 86 percent, or the vast majority of the variance.

In columns 2 and 3, we include product variety, measured by the (square root of) the firm count or by the barcode count, respectively.³² In either case, trade costs

³²The standard errors in column (2) are adjusted to reflect the error in estimating the (approximate) square root exponent that is applied to the firm counts (see note 27).

TABLE 2—DIFFERENCE BETWEEN THE COST OF LIVING AND THE TRADED CONSUMPTION PRICE LEVEL

Explanatory variable: $\ln(CoL^{Tr}/PC^{Tr})$	No variety (1)	Firm count (2)	Barcode count (3)
<i>Dependent variables:</i>			
<i>Trade costs and terms of trade</i>	0.106 (0.063)	0.079 (0.026)	0.090 (0.056)
<i>Sectoral share</i>	0.033 (0.313)	0.020 (0.015)	-0.005 (0.025)
<i>Inverse openness</i>	0.861 (0.080)	0.228 (0.066)	0.305 (0.139)
<i>Variety</i>		0.673 (0.045)	0.611 (0.106)
Number of countries	47	46	24

Notes: Each line in the table corresponds to a γ_k coefficient estimated from equation (20). Robust standard errors are in parentheses.

plus the terms of trade explain less than 10 percent of the variance in the cost of living relative to the consumption price level, and the sectoral shares account for only a slight (and insignificant) amount. Inverse openness now explains between 23 percent and 31 percent of the variance in the cost of living, while overall product variety explains between 62 percent and 68 percent. So product variety has overtaken inverse openness as the dominant explanation for how the cost of living across countries differs from the consumption price level.

IV. Results on Country Welfare

The above results have focused on the prices of tradable goods, for which we were able to collect measures of product variety. In this section, we extend those results by adding the prices of nontraded goods and by constructing a measure of welfare. We have not collected product variety or made any other adjustment to nontraded prices, but we simply combine them with our theoretical cost of living measure and with traded consumption price levels. By dividing consumption expenditure in each country by the theoretical cost of living, CoL , we obtain a theoretical measure of welfare, and if we use the price of consumption, PC , we obtain real consumption. These measures are compared to each other, and we compare our welfare measure to that from Jones and Klenow (2016).

A. Adding Nontraded Goods

Countries differ substantially in who pays for nontraded products. For health care and education, how much of these are purchased directly by consumers and how much by the government varies considerably across countries. Yet regardless of who pays for them—be it households, nonprofit organizations, or the government—these services are consumed. We thus use a measure of total consumption that includes

these services, corresponding to the statistical concept of actual individual consumption (*AIC*) for international comparisons.³³

We rely on ICP data to calculate the price level of nontraded consumption that is included in *AIC*, denoted by PC^{Ni} in country i . We add these nontraded prices into our previous calculation of the cost of living for traded goods consumption, to obtain

$$(21) \quad \frac{CoL^i}{CoL^j} \equiv \left(\frac{CoL^{Ti}}{CoL^{Tj}} \right)^{\omega^{Ti}} \left(\frac{PC^{Ni}}{PC^{Nj}} \right)^{\omega^{Ni}},$$

where the Sato-Vartia weights satisfy $\omega^{Ti} + \omega^{Ni} = 1$ (see online Appendix A.2). Likewise, the price level of *AIC*, inclusive of the nontraded services, is constructed as

$$(22) \quad \frac{PC^i}{PC^j} \equiv \left(\frac{PC^{Ti}}{PC^{Tj}} \right)^{\omega^{Ti}} \left(\frac{PC^{Ni}}{PC^{Nj}} \right)^{\omega^{Ni}}.$$

We are adding the same nontraded prices to both the theoretical cost of living and to the price of consumption goods, so that procedure will tend to reduce any differences in these two measures. Our final step is to use (21) or (22) to deflate nominal *AIC* in US dollars for each country relative to the United States, to obtain theory-based welfare ($U = AIC/CoL$), as compared to real consumption ($RC = AIC/PC$).

The results are shown in Figure 7, where, as in Figure 6, we used barcode counts to measure product variety for the 24 countries where these data are available (shown in blue), and otherwise the (square root of) firm counts for another 23 countries. In the left panel of Figure 7, we show the ratio of our theory-based welfare measured relative to real consumption. Because this figure is measuring welfare rather than the cost of living, it is roughly a mirror image of the right panel in Figure 6. There are some differences, however, because we have incorporated nontraded prices. Following countries from the left to the right, India (IND) and China appear first and have welfare from the Melitz model that is very close to real consumption. India had a higher cost of living relative to the United States than its traded consumption price level in Figure 6, but its high share of nontraded goods moves its welfare very close to real consumption in Figure 7. For the same reason, Indonesia (IDN) has welfare closer to real consumption in Figure 7 than we would expect from its relatively high theoretical cost of living in Figure 6. But aside from these differences, the set of countries with low theoretical cost of living in Figure 6 (Belgium, Denmark, Germany, Luxembourg, and the Netherlands) are still the countries for which welfare exceeds real consumption in Figure 7.

³³We only exclude net purchases of households abroad, which cannot be allocated to a type of products. In their macro-level comparison, Jones and Klenow (2016) focus on an even broader measure of consumption that also includes expenditure on collective goods and services. Collective services make up 10–12 percent of the Jones-Klenow consumption measure for most of our set of countries.

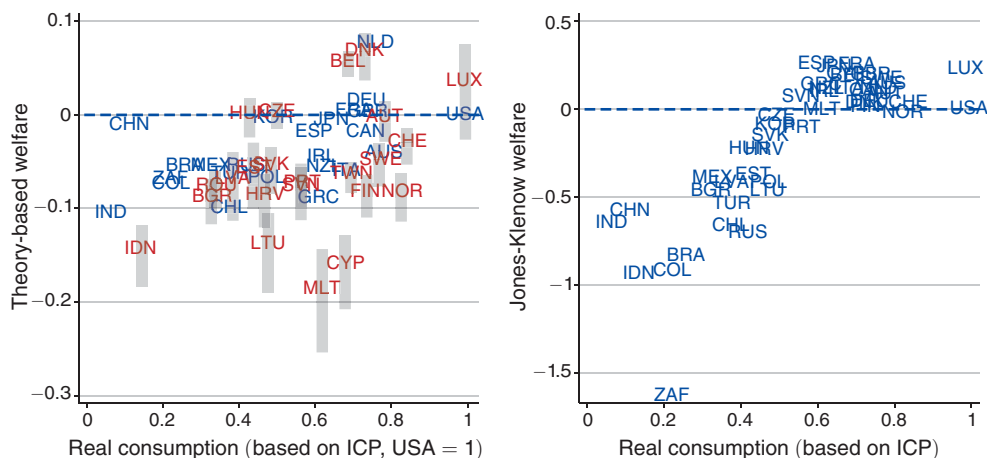


FIGURE 7. RATIO OF THEORY-BASED REAL CONSUMPTION TO REAL CONSUMPTION USING ICP PRICES, VERSUS REAL CONSUMPTION (LOG SCALE, USA = 0)

Notes: The left-hand figure plots the natural log of the ratio of theory-based welfare (using barcode counts for 24 countries, in blue, and firm counts for 20 countries, in red) to real consumption (based on ICP prices) against log real consumption, for 44 countries in our sample. Real consumption is computed by deflating actual individual consumption relative to the United States by the Sato-Vartia price index in (22), unlike the GEKS price index that is used by the ICP. The right-hand panel plots the log of ratio of welfare from Jones and Klenow (2016) to real consumption (based on ICP prices), against log real consumption, for the matching 44 countries in their sample.

B. Comparison with Jones and Klenow (2016)

Jones and Klenow (2016) propose a measure of welfare across nations that is meant to be much more inclusive than consumption, by also incorporating leisure, mortality, and inequality into a single consumption-equivalent measure. Our analysis, in contrast, is a more restrictive measure of welfare from the Melitz model that incorporates openness and product variety. Despite the differences in our approaches, it is worth asking whether the cross-country variation in welfare—as compared to a conventional measure of real consumption—has any similarity in their analysis and in ours. We find that they do.

The Jones-Klenow measure of welfare relative to real consumption for a matching set of countries as in our study is shown in the right panel of Figure 7.³⁴ The most obvious difference between the two panels is in the vertical scale of each: welfare in Jones and Klenow (2016) exceeds real consumption by up to 25 percent but can considerably lower as well. The most extreme difference is for South Africa (ZAF), where Jones-Klenow welfare for is only 4.5 percent of the US level, while real consumption

³⁴Not included are Hungary, Romania, and Taiwan, which were not covered by Jones and Klenow (2016). The direct correlation of our theory-based welfare measure and the Jones-Klenow measure (both relative to the United States) is 0.93 in levels and 0.83 in logs. One reason that these correlations are high, however, is that we have used ICP prices and actual individual consumption expenditures for nontraded goods both in our theory-based welfare and in real consumption, which is similar to the treatment of nontraded goods in PWT. Since Jones and Klenow (2016) also start with PWT data for real consumption (before adding leisure, mortality, and inequality), our measure of welfare and theirs will be correlated for that reason. Dividing both series by real actual individual consumption, as in Figure 7, eliminates that reason for correlation.

is 23 percent of the US level. Large differences are also seen for India (3.9 versus 7.4 percent) and China (6.6 versus 12 percent). Our measure of theory-based welfare differs by much less from real consumption, exceeding real consumption by up to 12 percent (for the Netherlands) and the largest reduction is 19 percent (in the case of Malta). The smaller scale in our case is not surprising in view of the more limited scope of our welfare measure.

Besides these differences, there are also similarities in the results. Many countries with low levels of real consumption have even lower levels of welfare, both in our results and in Jones and Klenow (2016). This is most pronounced in India (IND), Indonesia (IDN), Colombia (COL), and South Africa (ZAF), but also seen at higher consumption levels, such as for Chile (CHL) or Lithuania (LTU). As we move to the right, in either panel of Figure 7, welfare moves closer to real consumption, with only small differences in, for example, Czechia (CZE), Japan (JPN), and Korea (KOR).

The key difference between the two panels in Figure 7 is that welfare relative to the United States in nearly all the Western European countries *exceeds* real consumption for Jones and Klenow (2016), whereas in our case, welfare is higher only for Belgium, Denmark, Germany, Luxembourg, and the Netherlands.

V. Conclusions

The monopolistic competition model suggests that product variety is an important determinant of welfare. There are two challenges with evaluating this hypothesis. First, the most disaggregate data for measuring product variety—which are barcode data—are not typically available across multiple countries with the same classification system.³⁵ In the absence of a common classification system across many countries, we have relied on the count of barcode items from microdata in the BPP and, when those data are not available, on the simple count of firms as a proxy for variety.

Second, the literature on the gains from trade under monopolistic competition (ACR 2012) has emphasized the gains within a country when foreign variables, such as trade costs, change. To evaluate welfare between countries, however, we also need to include different domestic variables, such as domestic trade costs, productivity and fixed costs, and home population. We develop a parsimonious expression for the “true” costs of living in the Melitz model that incorporates changes to all these domestic variables and, therefore, endogenous changes in product variety. Because we compare the theoretical cost of living with the price level of consumption as measured from ICP data, we also end up comparing the openness of a country (which lowers the theoretical cost of living) with the terms of trade (which lower the price level of consumption relative to the price of output). Differences between openness and the terms of trade lead to commensurate differences between the cost of living and the consumption price level.

Before adjusting for product variety, more than half the countries in our sample have a cost of living from the Melitz model that is below their consumption

³⁵When it is available, as in Argente, Hsieh, and Lee (2022) for Mexico and the United States, and Beck and Jaravel (2020) for a broad set of countries, then it becomes possible to construct exact consumer price indexes as in Feenstra (1994). See also footnote 1.

price relative to the United States. Those differences are principally explained by the countries' openness as compared to their terms of trade. The United States, however, has higher product variety than nearly all other countries. Therefore, once we incorporate variety, the relative cost of living is raised in many countries, and we find that only five countries—Belgium, Denmark, Germany, Luxembourg, and the Netherlands—have costs of living relative to the United States that are below their consumption prices. A further group of European countries—including Austria, Czech Republic, France, Hungary, and the United Kingdom—along with China, Japan, and South Korea have costs of living that are similar to their consumption prices relative to the United States.

We have also used cost of living and consumption price levels to compute a theoretical measure of welfare and compare it with real actual individual consumption across countries while adding nontraded goods. Our theoretical measure of welfare varies inversely with the theoretical cost of living: Belgium, Denmark, Luxembourg, Germany, and the Netherlands have welfare relative to the United States above real consumption, while much the same group of European countries, along with China, India, Japan, and South Korea, have relative welfare that is not substantially different from real consumption; the remaining set of countries have lower relative welfare. That pattern is more pronounced in Jones and Klenow (2016), where nearly all Western European countries have welfare relative to the United States exceeding real consumption, whereas a smaller set of middle-income countries have lower welfare. It is surprising—but perhaps reassuring—that our narrow focus on the determinants of welfare in the Melitz model leads to a pattern of welfare across countries that has similarities to Jones and Klenow (2016) even though they focus on much broader determinants of welfare.

Our results raise the question of whether product variety should be incorporated into real GDP as computed by the ICP and PWT or kept in a separate account that is intended for research use. Our preference is for the latter approach. Just like the factors incorporated into welfare by Jones and Klenow (2016) are “beyond GDP,” we believe that the results for product variety are too preliminary and the country coverage too limited to be included now in official statistics. Still, we hope that with further research, the correction for product variety can become an accepted component of the cost of living and welfare across countries.

REFERENCES

- Argente, David O., Chang-Tai Hsieh, and Munseob Lee. 2022. “Measuring the Cost of Living in Mexico and the US.” *American Economic Journal: Macroeconomics* 15 (3): 43–63.
- Arkolakis, Costas. 2010. “Market Penetration Costs and the New Consumers Margin in International Trade.” *Journal of Political Economy* 118 (6): 1151–99.
- Arkolakis, Costas, Arnaud Costinot, and Andrés Rodríguez-Clare. 2012. “New Trade Models, Same Old Gains?” *American Economic Review* 102 (1): 94–130.
- Beck, Günter W., and Xavier Jaravel. 2020. “Prices and Global Inequality: New Evidence from Worldwide Scanner Data.” <https://ssrn.com/abstract=3671980>.
- Bureau Van Dijk. 2018. “Orbis Database.” Orbis. <https://orbis.bvdinfo.com> (accessed August 13, 2018).
- Cavallo, Alberto, W. Erwin Diewert, Robert C. Feenstra, Robert Inklaar, and Marcel P. Timmer. 2018. “Using Online Prices for Measuring Real Consumption across Countries.” *American Economic Review* 108: 483–7.

- Cavallo, Alberto, Robert C. Feenstra, and Robert Inklaar.** 2021. "Product Variety, the Cost of Living and Welfare across Countries." NBER Working Paper 28711.
- Cavallo, Alberto, Robert C. Feenstra, and Robert Inklaar.** 2023. "Data for: Product Variety, the Cost of Living, and Welfare across Countries." *American Economic Journal: Macroeconomics* 15 (4). <https://doi.org/10.34894/7RCSFZ>.
- Cavallo, Alberto, and Roberto Rigobon.** 2016. "The Billion Prices Project: Using Online Data for Measurement and Research." *Journal of Economic Perspectives* 30 (2): 151–78.
- Chaney, Thomas.** 2008. "Distorted Gravity: The Intensive and Extensive Margins of International Trade." *American Economic Review* 98 (4): 1707–21.
- Eaton, Jonathan, and Samuel Kortum.** 2002. "Technology, Geography and Trade." *Econometrica* 70 (5): 1741–79.
- Eaton, Jonathan, Samuel Kortum, and Francis Kramarz.** 2004. "Dissecting Trade: Firms, Industries, and Export Destinations." *American Economic Review* 94 (2): 150–54.
- Eaton, Jonathan, Samuel Kortum, and Francis Kramarz.** 2011. "An Anatomy of International Trade: Evidence from French Firms." *Econometrica* 79 (5): 1453–98.
- Feenstra, Robert C.** 1994. "New Product Varieties and the Measurement of International Prices." *American Economic Review* 84 (1): 157–77.
- Feenstra, Robert C., Robert Inklaar, and Marcel P. Timmer.** 2015. "The Next Generation of the Penn World Table." *American Economic Review* 105 (10): 3150–82.
- Feenstra, Robert C., Robert Inklaar, and Marcel P. Timmer.** 2021. "Penn World Table Version 10.0." Groningen Growth and Development Centre. <https://doi.org/10.15141/S5Q94M>.
- Feenstra, Robert C., and Hong Ma.** 2009. "Optimal Choice of Product Scope for Multiproduct Firms under Monopolistic Competition." In *The Organization of Firms in a Global Economy*, edited by Elhanan Helpman, Dalia Marin, and Thierry Verdier, 173–99. Cambridge, MA: Harvard University Press.
- Feenstra, Robert C., and John Romalis.** 2014. "International Prices and Endogenous Quality." *Quarterly Journal of Economics* 129 (2): 477–527.
- Feenstra, Robert C., Mingzhi Xu, and Alexis Antoniadis.** 2020. "What Is the Price of Tea in China? Goods Prices and Availability in Chinese Cities." *Economic Journal* 130 (632): 2438–67.
- Handbury, Jessie, and David E. Weinstein.** 2015. "Goods Prices and Availability in Cities." *Review of Economic Studies* 82 (1): 258–96.
- Hottman, Colin J., Stephen J. Redding, and David E. Weinstein.** 2016. "Quantifying the Sources of Firm Heterogeneity." *Quarterly Journal of Economics* 131 (3): 1291–1364.
- Jones, Charles I., and Peter J. Klenow.** 2016. "Beyond GDP? Welfare across Countries and Time." *American Economic Review* 106 (9): 2426–57.
- Kathuria, Sanjay, Arti Grover, Viviana Maria Eugenia Perego, Aaditya Mattoo, and Pritam Banerjee.** 2019. *Unleashing E-commerce for South Asian Integration*. Washington, DC: World Bank.
- Melitz, Marc J.** 2003. "The Impact of Trade on Intra-industry Reallocations and Aggregate Industry Productivity." *Econometrica* 71 (6): 1695–725.
- OECD-WTO.** 2018. "Trade in Value Added (Edition 2018)." OECD iLibrary. <https://doi.org/10.1787/ec7a5134-en>.
- Perla, Jesse, and Christopher Tonetti.** 2014. "Equilibrium Imitation and Growth." *Journal of Political Economy* 122 (1): 52–76.
- Perla, Jesse, Christopher Tonetti, and Michael E. Waugh.** 2021. "Equilibrium Technology Diffusion, Trade, and Growth." *American Economic Review* 111 (1): 73–128.
- Redding, Stephen J., and David Weinstein.** 2020. "Measuring Aggregate Price Indices with Taste Shocks: Theory and Evidence for CES Preferences." *Quarterly Journal of Economics* 135 (1): 503–60.
- Sampson, Thomas.** 2016. "Dynamic Selection: An Idea Flows Theory of Entry, Trade and Growth." *Quarterly Journal of Economics* 131 (1): 315–80.
- Simonovska, Ina and Michael E. Waugh.** 2014. "Trade Models, Trade Elasticities, and the Gains from Trade." NBER Working Paper 20495.
- UNCTAD.** 2017. *B2C E-commerce Index 2017*. UNCTAD Technical Notes on ICT for Development, 2017. Geneva, Switzerland: UNCTAD.
- WIOD.** 2021. "World Input-Output Database 2016 Release, 2000-2014." WIOD. <https://doi.org/10.34894/PJ2M1C>.
- World Bank.** 2020. *Purchasing Power Parities and the Size of World Economies: Results from the 2017 International Comparison Program*. Washington, DC: World Bank.