

Recovering Within-Country Inequality From Trade Data*

Dorothee Hillrichs[†]

Université catholique de Louvain

Gonzague Vannoorenberghe

Université catholique de Louvain

September 3, 2021

Abstract

This paper develops a novel method to estimate inequality within a country based on what it imports. If preferences are non-homothetic, rich and poor individuals in a country have different consumption profiles. Observing imports can thus inform us about the income distribution in a country. The global availability of trade data allows us to estimate inequality using the same transparent and comparable method for a large sample of countries over time. Compared to conventional data, we feature an especially good coverage of developing countries. We provide a number of robustness checks and cross-validation exercises to gauge the performance of our method.

JEL classification: F14, D63.

keywords: International trade, inequality, non-homothetic preferences.

*We would like to thank Manuel Oechslin for stimulating discussions that sparked our interest in this project. Our thanks also go to Lucas Avezum, Kristian Behrens, Antoine Berthou, Micael Castanheira, Malik Curuk, Katharina Erhardt, Mery Ferrando, Harry Huizinga, Udo Kreickemeier, Julien Martin, Florian Mayneris, Yasusada Murata, Gianmarco Ottaviano, Guzman Ourens, Mathieu Parenti, Louis Raes, Sjak Smulders, M. Scott Taylor, Burak Uras, Vincent Vicard as well as seminar participants at Bogazici University, Erasmus University of Rotterdam, Göttingen University, Marburg University, UQAM, and Tilburg University. Thanks also to attendants at the following conferences: DEGIT 2017, ETSG 2017, RIEF 2017, EEA 2019, ENTER 2019, ETSG 2019, GEP-CEPR 2019, RGS 2019. This work was supported by the Fonds Wetenschappelijk Onderzoek – Vlaanderen (FWO) and the Fonds de la Recherche Scientifique – FNRS under EOS Project No. 30784531 “Winners and losers from globalization and market integration” and under the CDR Grant J.0138.18 “Measuring inequality from trade data”.

[†]Place Montesquieu 3, 1348 Louvain-la-Neuve, Belgium, *e-mail:* dorothee.hillrichs@uclouvain.be

1 Introduction

Measuring income inequality within countries is a notoriously challenging task. Inequality data are typically based on surveys, on fiscal data or both. Surveys are however costly to conduct and fiscal data is available only where income is taxed. Both methods often lead to results that are difficult to compare across countries as survey questionnaires and fiscal rules are country-specific. Missing values are pervasive particularly among developing countries, where budget constraints and poor institutions make such data collection efforts daunting¹.

In this paper, we present a novel method for estimating inequality within a country that is based on what the country imports. Import data have a number of advantages over survey or fiscal data. They are easier to record than domestic activity, particularly for poor countries (Besley and Persson (2014)). They are double-recorded by the customs of the exporter and of the importer, and thus less sensitive to mismeasurement. They are widely and publicly available, and are recorded by all countries in a similar format, based on the “Harmonized System” of goods classification. We leverage these advantages to generate a comparable measure of inequality for virtually all countries.

Our method relies on the non-homotheticity of preferences, i.e. the fact that an individual’s income matters for his consumption basket. For example, consumers may switch from Greek, to Italian, and then to French wine when growing richer. Consider two countries with the same average income. The first country imports wine mostly from Italy while the second imports from France and Greece. Controlling for many other confounding factors (e.g. the distance from those countries, etc.), we take this as evidence of a higher inequality in the second country than in the first. Applying this logic over many imported products allows us to build a country-specific measure of inequality based on imports.

To guide our empirical analysis, we embed non-homothetic preferences in an otherwise standard model of international trade à la Armington, in which goods are differentiated by country of origin (e.g. French and Italian wines are two “vari-

¹For example, for the period 2016-2018, the United Nations’ World Income Inequality Database reports income inequality data for only 61 countries. Figure IX in the appendix breaks down the data coverage by region and income group of countries.

eties” of wine). This setup generates a non-homothetic gravity equation: alongside standard gravity determinants such as bilateral trade costs and multilateral resistance terms, the imports of a variety by a destination country depend on its average income and on the Gini coefficient of its income distribution².

Based on this structure, our empirical approach follows two steps. In the first step, we use a subsample of destinations with high-quality data on inequality and estimate for each variety a parameter capturing the extent to which it is imported by richer or more unequal countries. This parameter captures the relative income elasticity of a variety compared to the other varieties of a product. In terms of our example, this first stage allows us to rank the varieties of wines to which consumers turn as they grow richer. In the second step, we reverse the logic and estimate, for all countries, the Gini coefficient that best fits the observed import patterns given the variety-specific income elasticities estimated in the first step. Given a country’s average income, its distance from the exporter and other trade determinants, we predict the share of its product (e.g. wine) imports that should come from each exporter (e.g. France). We estimate the Gini coefficient that can best explain the deviations between the actual and the predicted imports over all varieties.

The validity of our method relies on a few key assumptions beyond the non-homotheticity of preferences. First, it requires some degree of comparability of preferences across countries. Using the sample of countries from our first step, we need to infer what people in other countries import when they grow richer. We require that, after controlling for standard gravity determinants of trade, if richer Germans consume relatively more French than Greek wine, it is also true of richer Brazilian consumers. It is worth emphasizing that we do not require that preferences for products are the same across countries. Some countries may have a strong preference for wine while others do not. What we need is that the relative income elasticities of varieties within products are the same. Second, since we rely on imports of consumer goods for identification, our logic requires that imports of goods provide enough information to reliably infer income inequality. If, for example, subsistence farmers consume few imported products or if rich individuals

²Recent research in international trade shows that the within-country income distribution plays an important role in explaining trade flows (e.g. Fielser (2011)).

fly to Paris to drink their wine, our inequality measure will miss some relevant variation. We discuss these issues at length and provide refinements to our method, based on a subsample of products for which such problems are less likely. We also discard country-time periods for which we deem our inequality estimates too sensitive to a few imported products.

Applying our method to 8 3-year periods from 1995 to 2018, our method generates a Gini coefficient for income inequality for more than 1000 country-period observations - almost twice as many as in the largest inequality data set available for that time period, the World Income Inequality Database. We add in particular observations on income inequality in Sub-Saharan Africa, the MENA region and South Asia. Our results reveal among other that the strong economic growth in South Asia in the last two decades was accompanied by rising income inequality, which approaches levels of Latin America. Such a comparison has been difficult to make with conventional survey data as these typically assess income inequality in Latin America but consumption inequality in South Asia.

We perform a number of out-of-sample comparisons to gauge the validity of the Gini coefficients and find that our estimates closely track changes in inequality measured by high-quality survey data. Where a point of comparison is available, we tend to find higher levels of income inequality than what is captured by survey data. For some countries, we diverge more strongly from existing sources. For Chile and Argentina, for example, our estimates show persistently high levels of inequality with only very modest improvements since the early 2000s while survey data point to an up to 10 Gini-point reduction in inequality. For the US, we find 5 to 10 Gini points higher inequality than survey data (depending on the comparison data source). These high levels are in line with values in the World Inequality Database that corrects survey data with administrative tax data.

Our approach differs substantially from most cross-country data collections on inequality, which we describe in the next section. We rely on comparable, public data that cover almost the whole world and apply the same, consistent method to all countries. Our method is however more indirect than using survey or fiscal data as it relies on observing imports patterns and not income directly. The cost of our

approach is the set of assumptions needed to map import patterns to inequality. We experiment with a number of robustness checks, modifying some of those assumptions, and present the results based on those alternative assumptions. The correlation of our inequality measures across robustness checks is very high. We see our paper thus as providing a simple and transparent method to complement the many careful studies providing a forensic analysis of particular countries by extending available data on inequality where surveys and fiscal data are scarce.

The remainder of the paper is structured as follows. Section 2 relates our method to other approaches taken in the literature on recovering inequality information, and gives an overview of the international trade literature our method is based on. Section 3 then outlines the theoretical structure we use to develop our method. Section 4 details the estimation strategy and discusses caveats to the method. Section 5 describes the trade and inequality survey data. Section 6 presents the estimated inequality data, section 7 validates our estimates with out-of-sample analyses and section 8 discusses the robustness of our method. Section 9 concludes.

2 Related literature

The purpose of this study is to provide consistently measured inequality data extending data coverage across countries and years. Therefore, we primarily contribute to the research on measuring inequality. We also add to the literature on non-homothetic preference in international trade.

2.1 Literature on inequality data

The problematic sparsity of inequality data based on surveys has gained the attention of researchers from various fields. McGregor et al. (2019) provide a comprehensive overview of challenges and solutions to measuring inequality. Ferreira et al. (2015) introduces a special issue of the *Journal of Economic Inequality* (Volume 13, Issue 4) devoted to assessing and comparing the most widely used inequality data sets available at the time. Two studies sharing the aim of extending the data coverage on inequality while ensuring consistency are undertaken by Galbraith and

Kum (2005) and Solt (2009). Galbraith and Kum (2005) predict household income inequality from a linear model of the Deininger and Squire (1996) Gini coefficient and the industrial payment inequality, controlling for the manufacturing share in employment as well as the Gini underlying welfare definition. Solt (2009) uses the information embedded in different types of inequality data to predict a standardized inequality index for countries lacking data. He argues for the Luxembourg Income Studies (LIS) as the inequality standard.

Both studies, Galbraith and Kum (2005) and Solt (2009), expand the coverage of inequality data. However, Solt has to rely on different methods for different countries to fill in missing data. Galbraith and Kum have difficulties in updating their data due to changes in the primary data used (Galbraith et al., 2014). Our method, in contrast, builds on economic theory and is applied uniformly to all countries using consistently collected trade data as the primary source of information. Changes in the classification of goods are well documented and thus do not pose a challenge to updating the data.

Another source of information on inequality is the World Inequality Database (WID) led by Facundo Alvaredo, Lucas Chancel, Thomas Piketty, Emmanuel Saez, and Gabriel Zucman. Their focus lies particularly on the upper tail of the income distribution. The researchers combine information from survey and fiscal data to derive a detailed account of a country's income distribution. While their forensic approach may lead to more accurate and nuanced inequality measures, data requirements are high and hard to be met in developing countries. Our approach complements their work in that it is applicable to all countries uniformly under comparatively low data requirements. As we show in section 6, our estimates are close to WID for many countries.

Recent advances in the field turn to secondary data, as we do, to uncover changes in the income distribution. Blumenstock et al. (2015), for example, map usage of mobile phones to the distribution of wealth in Rwanda. While a promising avenue for future research, this type of data is hard to obtain for any one country and far from lending itself to generate globally comparable data. Lessmann and Seidel (2017) use satellite nighttime light imagery to generate within-country regional

inequality data.

Our method is closely connected to Aguiar and Bils (2015). Aguiar and Bils (2015) measure consumption inequality for the US based on the relative allocation of spending of rich and poor households across luxuries and necessities. They employ a similar two-step estimation procedure to us. Their purpose, however, is different from ours. The goal of Aguiar and Bils (2015) is to address a systematic measurement error in the US Consumer Expenditure Survey. Instead, this study aims to generate a cross-country dataset on inequality.

The methodology of Almås (2012) is similar in spirit to our own. Almås (2012) estimates PPP corrected incomes across countries from estimating Engel curves for food. Similarly, our inequality measure exploits the variation in spending allocation across goods at different income levels. Our estimation, though, includes a larger variety of goods and is applied to a larger set of countries. The outcome variable of interest also differs between our study and Almås (2012).

2.2 Non-homothetic preferences in the trade literature

Our method builds on insights from the international trade literature. Inspired by the seminal dissertation of Linder (1961), a growing number of studies draw attention to the role of the income distribution as a determinant of trade patterns. Hallak (2006, 2010) provides evidence that similarity in the per capita income increases trade flows between countries. Choi et al. (2009) map income similarity to import price similarity suggesting a role for a country’s income distribution in explaining the quality of imported goods. Caron et al. (2014) show how import demand varies differentially with per capita income across goods. The authors estimate sector-level income elasticities which they use to assess the role of per capita income in explaining several trade “puzzles”. In this study, we provide further evidence on the significance of per capita income as a determinant of bilateral trade flows. We are the first to exploit the variation of income elasticities across goods to derive a measure of a country’s income distribution.

Theoretically, the role of the income distribution in determining trade patterns can be rationalized by non-homothetic preferences. The international trade liter-

ature suggests various microfoundations generating non-unitary income elasticities across goods. We introduce non-homotheticities similar to Faber and Fally (2017) and Handbury (2019) through a taste shifter in a CES framework. Alternative approaches include Comin et al. (2015) and Matsuyama (2019) who work with an implicit additively separable CES function. The resulting Engel curves are isomorphic to ours. Feenstra and Romalis (2014) employ the indirect utility of a CES function again implying very similar Engel curves. Fielor (2011) introduces non-homothetic preferences into a CRRA structure. In her model, the income elasticity of demand is governed by the elasticity of substitution. Relating to the consumer choice literature, Fajgelbaum et al. (2011) propose a nested logit demand system in which the expenditure share on quality varies with per capita income. Fajgelbaum and Khandelwal (2016) estimate heterogeneous income elasticities from an Almost Ideal Demand System.

Given the established high prediction power in log-linear models for trade flows, and the higher sensitivity to outliers in an estimation with import shares as dependent variable (instead of log imports, see Hillrichs and Vannoorenberghe (2021)), we opt for the class of non-homothetic CES preferences.³

3 Theory

We first present the approach in a general framework. To derive our estimating equation, we then put a parametric structure on the general model. For ease of exposition, we abstract from the time dimension and integrate it explicitly when describing our estimation strategy in section 4.

3.1 The general idea

The economy consists of C countries, indexed by $d \in \{1, \dots, C\}$ and J goods, indexed by $j \in \{1, \dots, J\}$. We think of goods as being differentiated by country of origin à la Armington, and assume that a variety is a good produced by a particular origin

³The structural change literature typically invoke Stone-Geary preferences (Kongsamut et al., 2001). However, under Stone-Geary preferences the income effect vanishes at high levels of income, which is undesirable for our purposes.

country, indexed by $o \in \{1, \dots, C\}$.

Within each country, we classify individuals into income categories, indexed by $h \in 1, \dots, H$, such that individuals in h have an income I_h . In country d , there are N_{dh} individuals in income category h , and a total of N_d individuals. We denote the average income of country d as $I_d = \sum_h N_{dh} I_h / N_d$ and define the expenditure that an individual in income category h in destination country d spends on the variety of good j coming from o as $x_{jo}(\mathcal{B}, \mathbf{A}_d, I_h)$, where \mathcal{B} is a matrix of unobserved variety-specific coefficients. \mathbf{A}_d is a matrix of destination-specific variables such as the vector of prices or of some taste parameters that can be observed or proxied. Aggregating over all individuals in country d gives:

$$X_{jod} = \sum_h N_{dh} x_{jo}(\mathcal{B}, \mathbf{A}_d, I_h). \quad (1)$$

Using a second order Taylor approximation of x_{jo} around the average income I_d and simplifying, the imports of variety jo in country d become (see appendix A for derivation):

$$\text{Log}(X_{jod}) \approx \text{Log}(N_d x_{jo}(\mathcal{B}, \mathbf{A}_d, I_d)) + \frac{1}{2} \frac{\partial^2 x_{jo}(\mathcal{B}, \mathbf{A}_d, I_d)}{\partial^2 I_d} \frac{I_d^2}{x_{jo}(\mathcal{B}, \mathbf{A}_d, I_d)} G_d, \quad (2)$$

where G_d is the squared coefficient of variation of income in country d . The first term captures the consumption of jo that would prevail in d if all individuals had the average income I_d . The second term captures the effect of income inequality on the consumption of jo and is non-zero as long as the the Engel curves are not linear.

We assume that data on the distribution of income are observable with great accuracy for some countries ($d \in \mathbb{O}$) while they are not for others ($d \in \mathbb{U}$). Our aim is to recover the distribution of income (G_d) for countries where it is unobserved, i.e. for $d \in \mathbb{U}$. Thanks to the wide availability of high-quality international trade data, we can observe the left hand side of equation (2) for virtually all imported varieties and destinations d . We thus propose a mapping from trade flows to a measure of inequality for countries $d \in \mathbb{U}$.

Our method has two steps. The first is to recover \mathcal{B} by estimating (2) for

$d \in \mathbb{O}$. In the second step, we use the estimated matrix $\hat{\mathcal{B}}$ to compute the predicted consumption of each variety if all households had the country’s average income. We then obtain an estimate of G_d for $d \in \mathbb{U}$ by regressing the difference between actual imports and predicted imports without inequality on the degree of convexity of the Engel curve⁴.

To illustrate the role of the convexity of the Engel curve, consider reallocating one dollar from the poor to the rich, an increase in inequality. The poor decrease their spending on the variety while the rich increase theirs (we think of a “normal” variety to simplify exposition). If the Engel curve is convex, the decrease in consumption by the poor is small compared to the increase in consumption by the rich and the total consumption of the variety rises. Keeping average income constant, relatively large imports of a good with a convex Engel curve are a sign of high inequality. The opposite holds if the Engel curve is concave.

A strong advantage of our approach is that trade flows are notoriously easier to measure than domestic activity in many countries. Trade flows are also double-recorded by the importer and the exporter, which makes them less subject to errors and less impacted by the poor administrative capacity of some countries. We thus proxy for inequality with a unified and transparent method for almost all countries using publicly available data.

The validity of our method naturally relies on a few key assumptions. First, it requires *non-homotheticity of preferences* for at least some varieties. With homothetic preferences, imports would not be affected by the income distribution (the Engel curves would be linear), implying that observing imports would not provide any information about the income distribution. Second, it relies on some degree of *comparability of preferences* between countries (\mathcal{B} is not specific to d). We do not require by any means that preferences are the same across the world (\mathbf{A}_d can contain preference parameters). What we need is that, by observing countries in \mathbb{O} , we can infer some patterns about what people in $d \in \mathbb{U}$ import when they grow richer.

Section 3.2 makes our assumptions on preferences explicit. Third, since we rely on

⁴Our strategy assumes that average income per capita is measured correctly for all countries. While measurement error in income per capita data certainly exists, the lack of inequality measures seems much more of a concern. We could in principle extend our method to recover both income per capita and income inequality for countries with low-quality data but leave it for future work.

imports of consumer goods for identification, this requires that *imports of goods provide sufficient information* about consumption patterns to extract information about the full income distribution. This may for example be an issue if subsistence farmers in Africa do not consume any imported goods or if, when growing very rich, consumers stop importing French wine but travel to France to drink it. We discuss these issues at length in section 8.

3.2 A non-homothetic gravity equation

We now make a number of functional form assumptions bring equation (2) to the data as a non-homothetic gravity equation.

Consumers in country d derive utility from consuming a bundle of differentiated goods, indexed by $j \in \{1, \dots, J\}$, and a homogeneous good, with a price normalized to one, indexed by 0. Utility is given by a Cobb-Douglas aggregator over the consumption of all goods C_j :

$$U_d = \prod_{j=0}^J C_j^{\rho_{dj}}, \quad \text{with:} \quad \sum_j \rho_{dj} = 1. \quad (3)$$

Each of the differentiated goods consists of varieties, differentiated by country of origin in an Armington fashion, such that, for $j \geq 1$:

$$C_j = \left[\sum_o \varphi_{jod}(C_0)^{\frac{1}{1-\gamma_j}} c_{jo}^{\frac{\gamma_j}{\gamma_j-1}} \right]^{\frac{\gamma_j-1}{\gamma_j}}, \quad (4)$$

where $\gamma_j < 0$ is equal to one minus the good-specific substitution elasticity between varieties. Following Faber and Fally (2017) or Handbury (2019), we introduce non-homotheticities in the preferences by making the demand shifter $\varphi_{jod}(C_0)$ depend on the consumption of the homogeneous good. We assume that:

$$\varphi_{jod}(C_0) = \frac{\alpha_{jod} C_0^{\beta_{jo}}}{\sum_{o'} \alpha_{jo'd} C_0^{\beta_{jo'}}}. \quad (5)$$

φ_{jod} is composed of a taste parameter α_{jod} , shared by all individuals in country d , and a power function of the individual's consumption of the homogeneous good C_0 , normalized by $\sum_{o'} \alpha_{jo'd} C_0^{\beta_{jo'}}$.

An individual with income I in country d maximizes his utility subject to the budget constraint⁵ $C_0 + \sum_{j=1}^N \sum_o p_{jod} c_{jo} = I$, where p_{jod} is the price of variety jo in country d . We solve the consumer's optimization problem formally in the appendix A and show that, under some additional conditions about the distribution of prices, $C_0 = \rho_{0d}I$. This implies that an individual with income I in country d spends

$$x_{jod}(I) = s_{jod}(I)\rho_{dj}I, \quad \text{where:} \quad s_{jod}(I) = \frac{\alpha_{jod}I^{\beta_{jo}}p_{jod}^{\gamma_j}}{\sum_{o'} \alpha_{jo'd}I^{\beta_{jo'}}p_{jo'd}^{\gamma_j}}, \quad (6)$$

on variety jo . $s_{jod}(I)$ denotes the share of this individual's spending on good j that is allocated to variety jo . The relative spending on a given variety within a good thus differs across individuals at different income levels.

The income elasticity of demand for variety jo for an individual in country d is:

$$\frac{\partial x_{jod}(I)}{\partial I} \frac{I}{x_{jod}} = 1 + \beta_{jo} - \bar{\beta}_{jd}(I), \quad (7)$$

where

$$\bar{\beta}_{jd}(I) = \sum_o s_{jod}(I)\beta_{jo}. \quad (8)$$

It can easily be seen that $\partial \bar{\beta}_{jd} / \partial I > 0$, i.e. richer individuals consume varieties with a higher average β . Some varieties jo , which may be inferior for some relatively rich individuals ($\beta_{jo} - \bar{\beta}_{jd}(I) < -1$), will however be normal goods ($-1 < \beta_{jo} - \bar{\beta}_{jd}(I) < 0$) or even luxury goods ($\beta_{jo} - \bar{\beta}_{jd}(I) > 0$) from the perspective of poorer individuals.

We show in the appendix that:

$$\left. \frac{\partial^2 x_{jod}(I)}{\partial^2 I} \frac{I^2}{x_{jod}} \right|_{I=I_d} = \beta_{jo} - \bar{\beta}_{jd}(I_d) + (\beta_{jo} - \bar{\beta}_{jd}(I_d))^2 - \sum_{o'} s_{jo'd}(I_d) (\beta_{jo'} - \bar{\beta}_{jd}(I_d))^2. \quad (9)$$

β_{jo} is closely related to the convexity of the Engel curve but in a non-monotonic manner. For $\beta_{jo} = \bar{\beta}_{jd}$, the Engel curve is concave, implying that a mean-preserving rise in income inequality reduces the consumption of the variety with the average β within a good. For varieties with a relatively high β_{jo} , the Engel curve is convex.

⁵We discuss deviations from this budget constraint in the robustness section, such as the possibility for consumers to save part of their income or for firms to charge different prices to consumers with different income even within a country.

Similarly, for varieties with a relatively low β_{jo} , the Engel curve is again convex as an increase in income inequality is associated with a stronger decrease in spending for the poor than for the rich.⁶ A mean-preserving rise in inequality thus also raises the consumption of varieties with a very low β_{jo} .

Each variety jo is produced under perfect competition, with constant marginal costs w_{jo} . Exporters of good j from country o have to ship $\tau_{jod} \geq 1$ units for one unit to arrive at destination d . These assumptions ensure that there is no pricing to market across countries ($p_{jod} = \tau_{jod}w_{jo}$), an assumption we come back to in section 8. In this case, the spending on variety jo by country d if all consumers had the average income - the first part of (2) - resembles a standard gravity equation, except for a coefficient on income per capita that is variety-specific:

$$\text{Log}(N_d x_{jod}(I_d)) = F_{jo} + F_{jd}^1 + \gamma_j \text{Log}(\tau_{od}) + \beta_{jo} \text{Log}(I_d) + \text{Log}(\alpha_{jod}) \quad (10)$$

where

$$F_{jd}^1 \equiv \text{Log}(N_d \rho_{jd}) - \text{Log} \left(\sum_o \alpha_{jod} I_d^{\beta_{jo}} p_{jod}^{\gamma_j} \right) + \text{Log}(I_d)$$

$$F_{jo} \equiv \gamma_j \text{Log}(w_{jo}).$$

The full equation (2) can thus be rewritten as:

$$\begin{aligned} \text{Log}(X_{jod}) = & F_{jo} + F_{jd} + \gamma_j \text{Log}(\tau_{od}) \\ & + \beta_{jo} \text{Log}(I_d) + \frac{1}{2} \beta_{jo} (1 + \beta_{jo} - 2\bar{\beta}_{jd}(I_d)) G_d \\ & + \text{Log}(\alpha_{jod}) \end{aligned} \quad (11)$$

where:

$$F_{jd} \equiv F_{jd}^1 + (2(\bar{\beta}_{jd}(I_d))^2 - \sum_o s_{jod}(I_d) \beta_{jo}^2 - \bar{\beta}_{jd}(I_d)) G_d. \quad (12)$$

⁶Note that, from equation (9), a necessary condition for the Engel curve to be convex when $\beta_{jo} - \bar{\beta}_{jd} < 0$ is that the good be inferior, i.e. $\beta_{jo} - \bar{\beta}_{jd} < -1$.

4 Estimation strategy

4.1 A two-stage approach

The first step of our approach is to identify the coefficients β_{jo} in equation (11) based on a subsample of countries for which we have reliable inequality data (\mathbb{O}). In a second step, we invert the relationship and use the estimated β_{jo} from the first step to recover inequality for countries in \mathbb{U} . We will use data for different periods and now introduce the time subscript explicitly.

4.1.1 First stage

In the first stage, we estimate (11) product by product using Non-Linear Least Squares (NLS). We control for F_{jot} using an exporter-product-time fixed effect and for F_{jdt} using a destination-product-time fixed effect. We capture bilateral trade costs τ_{od} using geographical distance, dummies for shared border, official language and colonial past as well as participation in the same regional trade agreement. We use real GNI per capita to capture the average income in the destination (I_{dt}) and interact its logarithm with a dummy variable that takes value one for trade flows of variety jo . This allows the coefficient on $\text{Log}(I_{dt})$ to be variety-specific. In a similar manner, we interact the squared coefficient of variation G_{dt} with a dummy variable for each variety to capture the term $\beta_{jo}(1 + \beta_{jo})\frac{1}{2}G_{dt}$. The NLS estimation guarantees that the relationship between the variety-specific coefficients on $\text{Log}(I_{dt})$ and G_{dt} is in line with our model. To capture the term $\beta_{jo}\bar{\beta}_{jd}\frac{1}{2}G_{dt}$, we group countries according to their patterns of imports. We form 3 groups of countries per product⁷ based on the similarity of their vector s_{jod} . In line with (8), countries with similar import patterns should have a similar value of $\bar{\beta}_{jd}$ and we thus assume $\bar{\beta}_{jd} = \bar{\beta}_{jg(d)}$, i.e. that it is common to all countries of the group in which d is $g(d)$. We introduce an interaction between G_{dt} , a dummy for the variety jo and a dummy for the group to which destination d belongs for product j . Finally, the error term ε_{odjt} captures unobservable taste shocks ($\text{Log}(\alpha_{jodt})$). To

⁷The details of the procedure are available in appendix section C. The groups are specific to HS chapters.

summarize, our first-stage estimation relies on an NLS estimation of the following equation per product j :

$$\begin{aligned}
\text{Log}(X_{jodt}) &= F_{jot} + F_{jdt} + \gamma_j \text{Log}(\tau_{od}) \\
&+ \sum_{o'} \beta_{jo'} \mathbb{1}_{o'=o} \text{Log}(I_{dt}) \\
&+ \sum_{o'} (\beta_{jo'} + \beta_{jo'}^2) \mathbb{1}_{o'=o} \left(\frac{1}{2} G_{dt} \right) + \sum_{o'} \sum_{g'_j} b_{jo'g'_j} \mathbb{1}_{o'=o} \mathbb{1}_{d \in g'_j} \left(\frac{1}{2} G_{dt} \right) \\
&+ \varepsilon_{jodt},
\end{aligned} \tag{13}$$

where $\mathbb{1}_{o'=o}$ is the indicator function that takes value one for exports from origin o and zero for all other varieties, and where $\mathbb{1}_{d \in g'_j}$ takes value one for all countries importing variety jo that are in group g'_j .

4.1.2 Second stage

In the second stage, we exploit the variation of income elasticities across varieties to infer changes in a country's income distribution from changes in import patterns.

For this, we define:

$$\widetilde{\ln(X_{jodt})} \equiv \ln(X_{jodt}) - \hat{F}_{jot} - \hat{\gamma}_j \ln(\tau_{od}) - \hat{\beta}_{jo} \ln(I_{dt}), \tag{14}$$

where $\hat{\gamma}_j$, $\hat{\beta}_{jo}$ and \hat{F}_{jot} are our parameter estimates from the first stage. According to our model, variations in $\widetilde{\ln(X_{jodt})}$ should only be explained by destination-product-time fixed effects, idiosyncratic taste shocks, and inequality in the distribution of income. Given the curvature of the Engel curve, larger deviations of $\widetilde{\ln(X_{jodt})}$ from observed imports imply a larger disparity of spending on variety jo across households of different income. Our strategy is thus to regress $\widetilde{\ln(X_{jodt})}$ on the estimated inequality (semi-)elasticity $\hat{\beta}_{jo}(1 + \hat{\beta}_{jo} - 2\hat{b}_{jog_j(d)})$ interacted with country-time dummies (D_{dt}), controlling for product-destination-time fixed effects⁸. The

⁸The second step considers country-time pairs that were not part of the first stage. We therefore do not have a first stage estimate for these fixed effects.

second step regression equation is:

$$\widetilde{\ln(X_{odjt})} = f_{jdt} + D_{dt}(\hat{\beta}_{jo}(1 + \hat{\beta}_{jo} - 2\hat{b}_{jog_j(d)})) + \nu_{jodt}. \quad (15)$$

The coefficients on the interaction term are the estimates of the squared coefficient of variation for each country and time period. Essentially, we switch what is a variable and what is a parameter in equation (13) when moving from step one to step two. The inequality estimates are identified from the variation of income elasticities (determined by $\hat{\beta}_{jo}$) across varieties. Due to the presence of the fixed effect, it is the within-product variation across varieties on which the identification relies.

4.2 Identification

As explained in section 3.1, our method relies on some degree of comparability of preferences across countries. This does however not mean that all preference parameters need to be equal in all countries. For example, the Cobb-Douglas weights on good j (i.e. the total spending on j , or ρ_{jd}) can differ across destinations and are captured by destination-good-time fixed effects in our estimation. Our identification strategy precisely aims to limit issues related to preference comparability by using *within-product* variation between varieties. Our key assumption in terms of comparability of preferences is that β_{jo} is common to all destination countries. This implies that, when they grow richer (or poorer), consumers anywhere tend to turn to the same varieties of a good.

We also allow for countries to differ in terms of a destination-variety specific preference shifter, our structural error term. In the first stage, a consistent estimate of β_{jo} for a variety requires that the error term be uncorrelated with the income or inequality of destination countries. We think that this assumption will likely hold for many varieties, but it might sometimes be violated⁹. For example, if Swiss watches are particularly popular in countries that host a Grand Slam tennis

⁹We could in principle use more fixed effects, for example at the product-origin-destination level but while theoretically possible, relying only on within-country variation would drastically reduce the sources of variation.

tournament and these countries are richer than average, this may drive an upward bias in the β of Swiss watches. A biased first stage estimate would however only pose a threat to our strategy if it carries over to the second stage.

In the second stage, the existence of an idiosyncratic preference shifter biases the estimate of the Gini coefficient if it is correlated with the estimated income elasticity across varieties. In other words, it would require that a country disproportionately likes varieties typically preferred by the world's rich (or poor) for reasons that have nothing to do with income or gravity variables. If the estimated β_{jo} in the first stage is biased for some varieties, another issue would be a correlation between those first-stage biases and the country's preference shifters across all varieties. We see such correlations as quite unlikely over the hundreds of products that we consider. To be conservative, we present our results about the Gini coefficient both in levels and in changes. Any time-invariant correlation between the preference shifters and the estimated income elasticity across varieties would not affect the evolution of the Gini coefficient over time, even if it may bias its level.

5 Data

5.1 Trade data

We choose the BACI trade database compiled by CEPII, which is based on the UN Comtrade database, to obtain bilateral trade data at the product level. The BACI trade database reconciles trade flow reports by the importer and the exporter country (for details of the methodology see Gaulier and Zignago (2010)). This is important for our purposes. Exploiting the double-recording nature of trade flows lessens the dependency on the statistical capacity of a single country regarding the data quality. Products are classified by the Harmonized System (HS). We aggregate trade flows to the HS 4-digit product level to limit the incidence of zero trade flows and maintain a sufficient number of varieties by product. We retain only consumption goods, as classified by UNCTAD, since expenditure on these goods should be relatively more affected by per capita income and inequality than raw materials or intermediate goods.

We observe trade flows for the period 1995 to 2018. To smooth out annual shocks, we aggregate the trade flows to three-year periods and take the three-year average value for yearly varying control variables. To increase the reliability of our data further, we exclude trade flows to or from countries with fewer than 500 000 inhabitants. We keep only the largest exporters, the top half of the export value distribution within a product over the entire sample period. Finally, we restrict the sample to exporters with more than 25 destination-time observations per good as well as importers with more than 25 imported varieties per period to maintain sufficient estimation precision. This latter restriction is binding for less than 1% of our data.

Table I presents a summary of our trade data. Our final sample consists of 365 HS 4-digit codes, 122 exporting countries, and 161 importing countries. In the first stage, variety-specific income elasticities are identified from the variation of per capita income and net income inequality across on average 328 country-periods of our first step sample. The number of country-periods per variety in the first step sample varies between 28 (Morocco, Articles of natural cork) and 446 (India, Medicaments for retail sale). The identification of within-country inequality in the second stage makes use of variation across on average more than 3800 varieties imported per period. The number of varieties imported per product and period ranges from 5 to 38.

Table I: Trade data, key statistics

	Mean	Sd	Median	Min	Max
Importer-time per variety (1st sample)	328	97	353	28	446
Varieties per importer & time & HS4	14	7	13	5	38
Varieties per importer & time	3,894	1,755	3,940	142	6,662

Notes: The table presents summary statistics on the number of observations along key dimensions in the trade data. In total, we have 7190 varieties in the first stage sample, 322360 unique importer-time-HS4 combinations, and there are 1230 unique importer-periods.

5.2 Within-country inequality data

The second key data used in our estimation is data on inequality. The measure of inequality that enters our estimating equation is the squared coefficient of variation of

a country’s income distribution (G), which is not usually available in cross-country databases of inequality. To connect theory to data, we follow Fajgelbaum and Khandelwal (2016) and rely on the following approximation

$$G_{dt} = \exp(4(\operatorname{erf}^{-1}(\operatorname{Gini}_{dt}))^2) - 1, \quad (16)$$

which holds under a log-normal distribution of income and where erf^{-1} is the inverse error function.

We obtain the within-country Gini data from the World Income Inequality Database (WIID Version 2020) administered by UNU-WIDER. The WIID is itself a collection of Gini data sets coming from various institutional surveys and research studies. Each data point is described by the underlying welfare concept (net/gross income, consumption), the reference population (urban/rural/all) and equivalence unit of analysis (person/household) and is categorized by a quality level (low/average/high), which refers to the transparency of the primary data collection. We set these categories such that we extract the largest data series with a uniform interpretation of the Gini coefficient.

We restrict our sample $dt \in \mathbb{O}$ from those country-year pairs for which the Gini in terms of disposable income is recorded. The choice of net income as welfare concept is consistent with our theory and furthermore motivated by the higher number of high quality Gini data in the WIID (Version 2020) compared to Gini data in terms of consumption, which would be a natural alternative welfare concept.¹⁰ This choice also dictates the interpretation of our trade-based Gini estimates as capturing net income inequality. We impose that the data source be ranked high quality, which is the case for most entries of net income Gini.¹¹ In addition, we set the unit of analysis to “person” and the reference population to “all”. Where the WIID still has duplicate entries for a country-period after applying our selection criteria we give preference to data from specific sources in the following order: the Luxembourg Income Study (LIS) 2019, OECD 2019, OECD 2018, ECLAC 2019,

¹⁰We apply our methodology also to the case of consumption inequality as a robustness exercise.

¹¹Only for the few country-periods for which WIID records no high quality entries, we include data from average quality sources if all other selection criteria are met.

Eurostat 2019. Table XV in the appendix gives a detailed list of all original Gini data sources.

In total, we include 443 destination-periods for our first stage estimation. A comparison of columns 1 and 2 in Table II shows that our selection lowers the number of country-period pairs with net income Gini data available in the WIID by 88. The map in Figure I presents the (importer) countries in our first step sample, i.e. those countries for which we observe the Gini for some period. Europe, North and Latin America is mostly well covered. Yet, data is scarce for African and many Asian countries.

The third column in Table II gives the number of countries for which we will be able to estimate a Gini index with our approach. Our method achieves a substantially larger coverage than the WIID. This is true even if one counts all entries in the WIID irrespective of the Gini definition and the quality rating for the same period (974 unique country-period entries).¹²

The final three columns present the coverage of countries in three data sets that we do not use in our estimation. These data sets serve as out-of-sample benchmarks to evaluate our Gini estimates in section 7. The first is an updated version of the LIS (version 2020 whereas WIID contains version 2019), the second is the World Bank's Povcal (mostly rated average quality by WIID and therefore not part of our first stage sample), and the third is the World Inequality Database (WID)¹³.

¹²The numbers presented in column 3 conceal however a certain variation in the stability and accuracy of our trade-based inequality estimates, an aspect under review at the time of submission to the Second World Inequality Conference

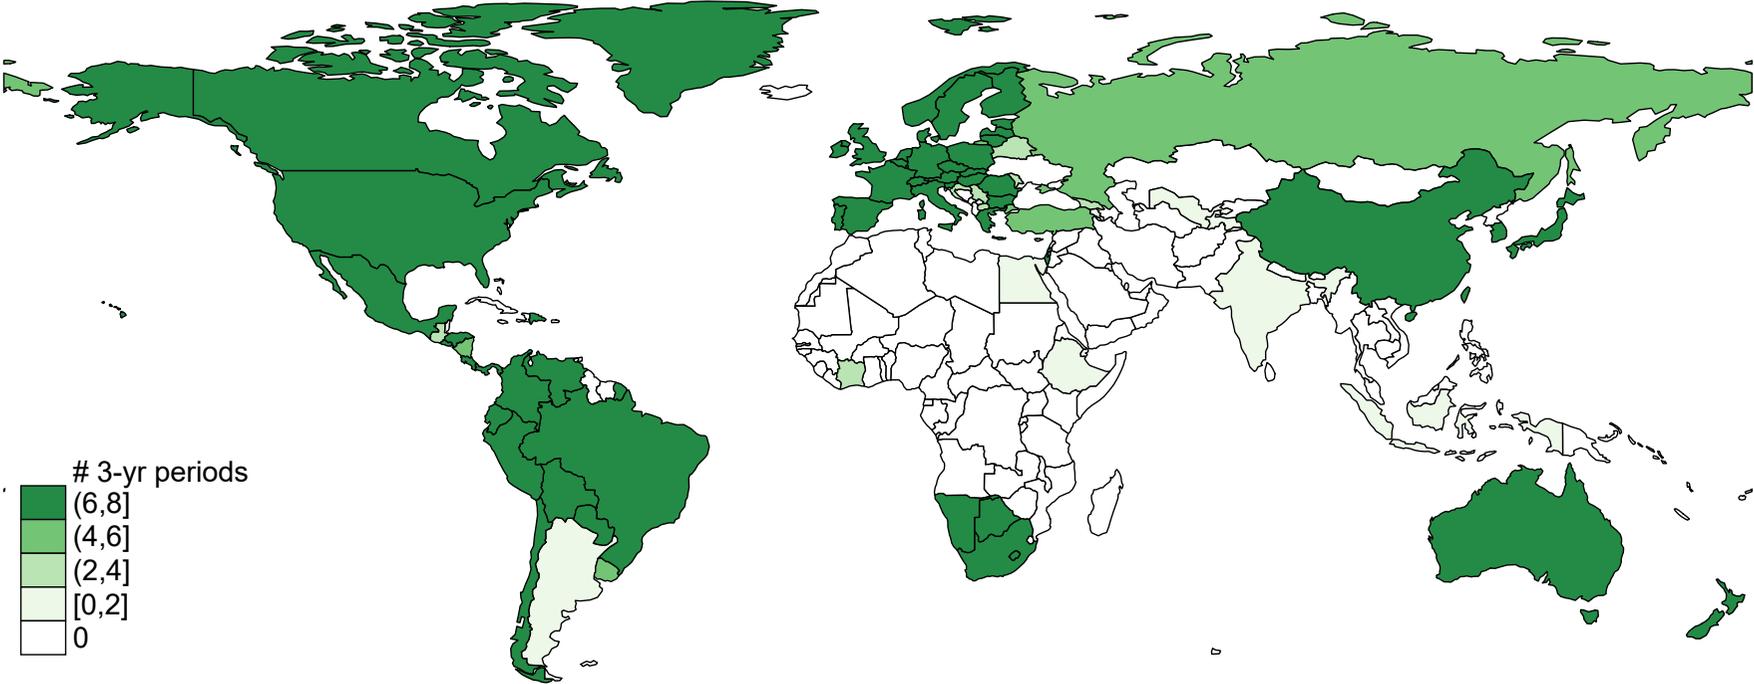
¹³Research at the World Inequality Database puts a stronger emphasis on recovering the income at the top percentiles of the distribution rather than the Gini index. The WID has more countries covered with information on the top part of the income distribution than for which it reports the Gini index.

Table II: Number of countries covered in various Gini datasets

Years	WIID 2020		trade- based estimate	LIS 2021	Povcal 2021	WID 2021
	all net income	selection				
1995-1997	62	47	141	25	36	37
1998-2000	67	54	146	27	40	37
2001-2003	73	53	149	13	38	37
2004-2006	69	57	148	40	55	38
2007-2009	64	56	153	41	56	38
2010-2012	68	61	154	46	57	38
2013-2015	67	60	154	43	60	38
2016-2018	61	55	153	35	51	38
Total	531	443	1198	270	393	301

Notes: Luxembourg Income Study Version March 2021. World Development Indicators accessed March 2021. World Inequality Database: accessed March 2021. WIID: Version May 2020. Gini definition: Net income inequality. Other selection criteria see appendix.

Figure I: Gini availability: Step 1 sample



Notes: The maps show the number of time periods for which within-country inequality is recorded in the WIID and fulfills our selection criteria (see text).

5.3 Country groups

To estimate the parameter β_{jog} , we design an algorithm to form groups of countries such that countries $dt \in \mathbb{O}$ (colored in Figure I) reflect the preferences of the countries $dt \in \mathbb{U}$ (blank in Figure I). The grouping is executed per HS Chapter (in total 15) and works as follows. We calculate an exporter's share in the 24-year aggregate expenditure for each importer. Next, we calculate the cosine similarity of importers in terms of these shares. We group the countries $dt \in \mathbb{O}$ into three groups using Ward's linkage method. The coefficient β_{jog} is identified from the variation in inequality across countries within such a group. For each country $dt \in \mathbb{U}$, we calculate the average similarity across countries $dt \in \mathbb{O}$ by group and assign the country $dt \in \mathbb{U}$ to the group with the highest average similarity across members. The appendix C describes the process in more details and presents the group composition by HS Chapter.

5.4 Other data

Our proxies for trade costs (bilateral distance, dummies for shared border, official language and colonial past as well as participation in the same regional trade agreement) come from CEPII and Egger and Larch (2008). We use real GNI per capita data provided by the World Bank. Lastly, for a robustness exercise we incorporate unit value data obtained from CEPII. Table 14 in the appendix provides a list of all variables we use including data sources.

6 Results

We first briefly discuss the results of the first step, the income elasticity estimates, before presenting our trade-based inequality measure. In the appendix, we provide our within-country income inequality estimates for all countries and all periods.¹⁴

¹⁴We also work on bootstrapped standard errors at the time of submission. We draw 100 random samples with replacement from our original data set for each exporter-product-sample, where sample refers to samples \mathbb{O} and \mathbb{U} (whether or not we observe the Gini coefficient of a country). We then apply the full procedure on each sample, apply equation (16) on the second stage estimates and compute standard errors from the 100 Gini estimates.

6.1 First stage: the income elasticities

In our first stage, we estimate separate gravity coefficients for each HS 4 code and the elasticities β_{jo} and β_{jog} for each of our more than 6000 varieties. Table III summarizes the distribution of our parameters. Seven products with extreme values for $\hat{\beta}_{jog}$ are omitted from the table, as we also drop them from the second stage estimation.¹⁵ The income elasticity estimates show a large heterogeneity within and across products. This is a necessity for our method as our second step relies on precisely this source of variation. The coefficients on the various trade cost proxies reported at the bottom of Table III are in line with conventional results.

To get a sense of the trade patterns underlying our income elasticity estimates, we plot in Figure II the weighted average income elasticity of an *exporter* against its log per capita income. Each variety's standardized¹⁶ income elasticity is weighted by the product's share in total exports of the country. Previous contributions to the trade literature (Hallak, 2006; Fajgelbaum and Khandelwal, 2016) suggest a positive link of trading countries' per capita income and hence would imply a positive correlation of an exporter's income and its associated average income elasticity. Instead, Figure II exhibits a U-shaped correlation. The negative correlation we find among exporters of low per capita income appears to be driven to a large degree by the textile industry in South-East Asian economies, which is highly integrated into global value chains. As firms from high-income markets optimize production costs and outsource labor-intensive, final production stages to these countries, it is not unreasonable to find large exports of final consumption goods from low-income to high-income countries. Unfortunately, value added data are not available at a comparable product disaggregation nor for a comparable number of countries as are data on gross import flows. For our method, integration into global value chains creates a challenge only if it implies that country exports different varieties to different markets depending on the destination's income. We address this concern in section 8.

¹⁵These are 7 products with $\hat{\beta}_{jog} > 10$ for some *og*.

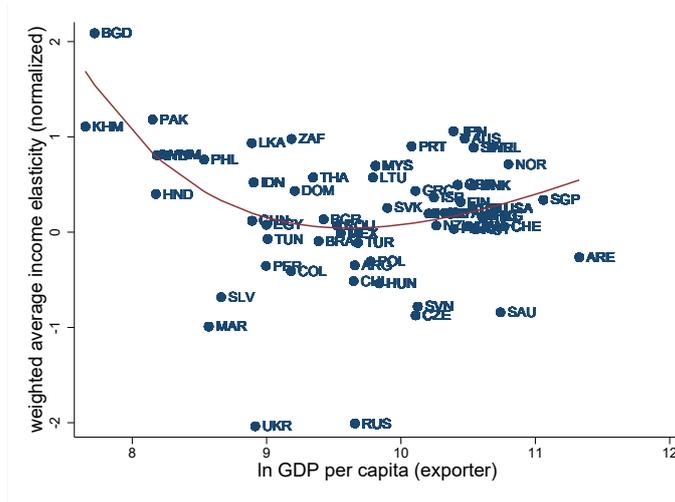
¹⁶Income elasticities are normalized by the product mean and standard deviation.

Table III: Summary Statistics of first stage results

	Mean	Sd	Median	Min	Max
$\beta_{jo}^{p75} - \beta_{jo}^{p25}$	0.78	0.28	0.73	0.18	2.09
distance	-1.02	0.32	-0.97	-1.84	-0.12
border	0.58	0.37	0.52	-0.53	2.51
language	0.42	0.28	0.38	-0.76	1.61
colony	0.52	0.29	0.56	-1.27	1.39
RTA	0.39	0.23	0.39	-0.33	1.78

Notes: Summary statistics of the distribution of coefficient estimates from equation (13) across 358 products. Row one reports the cross-product summary statistics of the difference in income elasticities at the 75th and 25th percentile of the within-product distribution of income elasticities.

Figure II: Income elasticity estimates across exporters



Notes: The figure shows the relationship between the weighted average of the normalized income elasticity $\frac{\hat{\beta}_{oj} - \tilde{\beta}_j}{s.d._j(\hat{\beta}_{oj})}$ and the natural logarithm of the exporter's average GDP per capita over the sample period 1995 - 2018. The weights are the product share in total exports over the sample period. Countries exporting less than 10 products are omitted from the graph.

6.2 Second stage: estimating inequality

Our second stage provides estimates of \tilde{G}_{dt} , which we transform into a Gini index using equation (16). We obtain a Gini index for net income, bounded between 0 and 100, with higher values indicating higher inequality.

Table IV presents summary statistics of our Gini estimates over the full sample period. Our method generates estimates for more than 1200 country-periods. We point out that with a minimum value of 9 and the first percentile of the Gini

distribution at 16 some values are however unlikely low. In the remainder of the paper, we drop estimates we deem to unstable in a bootstrap-like exercise.¹⁷ This leaves us with in total 1050 Gini observations.

Table IV: Summary statistics of Gini estimates

	Mean	Sd	Median	Min	1st perc.	Max	Obs
$Gini^{trade}$ (est.)	41.5	9.7	41.8	8.9	15.8	69.4	1,198

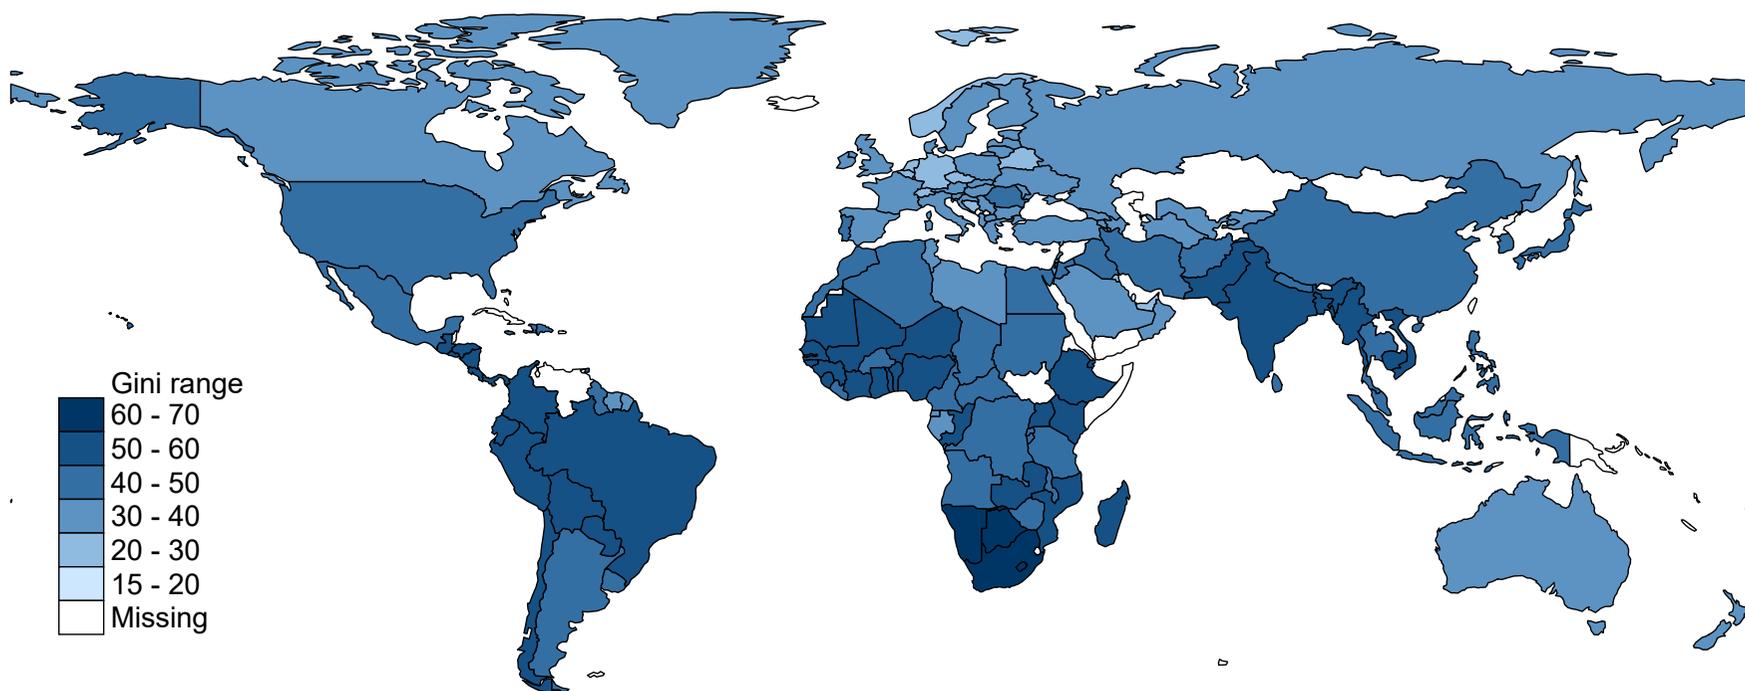
Notes: The table presents summary statistics of the Gini distribution based on the estimation of equation (15).

The world map of income inequality in Figure III depicts our trade-based inequality estimates for 2016-2018. Darker shades indicate higher inequality, countries in white lack data or are set to missing if the estimates are too unstable.¹⁸ As evident from the map, our method can shed some light on inequality in regions that would otherwise be left blank. We provide the full set of our Gini estimates in the appendix and on our research website.

¹⁷To give a more systematic assessment of the reliability, we are running a panel bootstrap on the second step using the empirical distribution of the data. We resample 100 times over HS 4-products for each country-period pair. This method is also referred to as block bootstrap. Here the “block” consists of the complete set of varieties of an HS-4 product that the destination imports at time t . The bootstrap results reveal how sensitive a Gini estimate for a given country-period is to a change in the set of products a country imports in that period. The method is being refined at the time of submission.

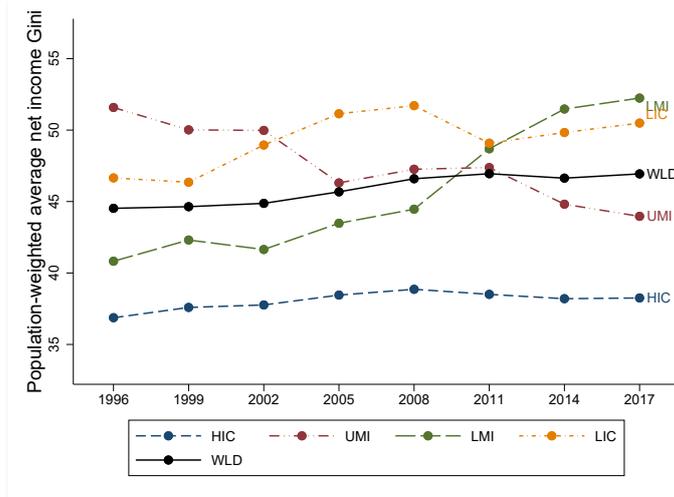
¹⁸Data may be missing for three main reasons: no trade data (e.g. North Korea), no data on GNI per capita (e.g. Venezuela 2016-2018) or a negative estimated G_{dt} . The latter happens for few time periods for United Arab Emirates, Kuwait, Macao, Laos, Bahrain, Oman, and Saudi Arabia).

Figure III: Net income inequality estimates, 2016-2018



Notes: The map show the within-country inequality for the period 2016-2018. Darker shades indicate higher income inequality. Within-country income inequality is estimated using equation (15) and transformed to the Gini index using equation (16).

Figure IV: Inequality 1995 - 2018 by country income group

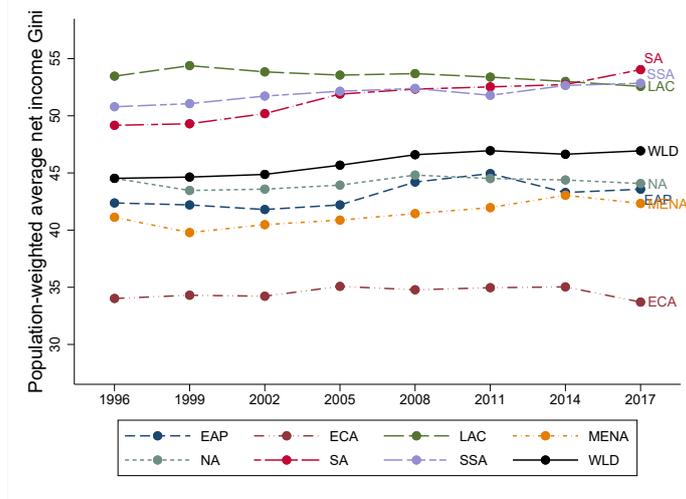


Notes: Population-weighted average Gini estimate by income groups: HIC - high income, UMI - upper middle income, LMI - lower middle income, LIC - low income, WLD - World average. Labels on the horizontal axis indicate the middle year of a three-year period.

To illustrate the contribution of our approach, we present here some descriptive graphs that would not be very meaningful without consistent data on inequality in large parts of Africa and Asia. First, Figure IV plots the evolution of inequality by income groups of countries (according to the World Bank definition). The population-weighted average within-country inequality has been rising between 1995 and 2018. Inequality is lowest among the most advanced economies. Notably lower-middle income countries experienced surging levels of inequality while inequality in upper-middle income countries appears to have been declining throughout.

Figure V plots the development of population-weighted average inequality over the years 1995 - 2018 by world region. A striking pattern of three inequality levels emerges. In line with survey data, Latin America stands out as one of the most unequal regions. At similarly high levels is inequality in Sub-Saharan Africa and South Asia. Whereas inequality was in decline in Latin America, economic growth of South Asian economies in the last two decades, notably India, appears to have been accompanied by the sharpest increase in inequality among all world regions. North America, the East Asian and Pacific countries and finally the Middle East & North Africa (MENA) region form the middle group. Our estimates show that,

Figure V: Inequality 1995 - 2018 by world region



Notes: Population-weighted average Gini estimate by regions: EAP - East Asia & Pacific, ECA - Europe & Central Asia, LAC - Latin America & Caribbean, MENA - Middle East & North Africa, NA - North America, SA - South Asia, SSA - Sub-Saharan Africa, WLD - World average. Labels on the horizontal axis indicate the middle year of a three-year period.

from the late 1990s onwards, inequality rose continuously for 15 years in the MENA region. Such development was not visible in the scarce survey data for countries in this region and as a result the Arab Spring revolts caught many observers by surprise.¹⁹ On the other side of the inequality spectrum, Europe and Central Asia have the lowest levels of inequality.

7 Validation of inequality estimates

We now turn to comparing our inequality estimates with other existing sources of data on inequality.

7.1 Comparison with WIID (out-of-sample)

One way of evaluating our method's performance is to use a k-fold cross-validation so as to avoid mechanical correlation between our Gini estimates and the WIID data we used in the first stage. We randomly split our first-step sample of country-time pairs with observed Gini data into ten subsamples. We drop each subsample in turn from the estimation of income elasticities, apply our two-step approach to

¹⁹The World Bank refers to this as the 'Arab inequality puzzle' in Ianchovichina et al. (2015).

Table V: 10-fold cross-validation results

	First-differences			Level		
	Mean	Sd	Median	Mean	Sd	Median
Baseline - Net Income Gini	2.43	2.24	1.85	4.64	3.28	4.15

Notes: The table presents the distribution of the absolute difference between the out-of-sample predicted Gini change/level and the change in WIID Gini (level of WIID Gini).

the remaining countries, and evaluate the Gini estimates against the data for those countries that were dropped from the first step.

Table V shows the absolute deviations of the out-of-sample predicted Gini from the first step Gini sample in terms of first differences and in terms of levels. On average, the difference in predicted and observed period-on-period change is approximately 2.4 Gini points, the median is 1.9 Gini points. In levels, the difference between out-of-sample estimate and observed Gini is, on average, 4.6 points (median: 4.15 points). No strong bias emerges either: roughly two-thirds of our estimates exceed the first stage Gini values, one third of the estimates lie below the WIID values.

7.2 Comparison with other data sets

While it is difficult to obtain comparable data on inequality for a large set of countries, many studies exist that measure inequality at the level of individual countries, for example using surveys. As alternative points of comparison for our estimates, we use the Luxembourg Income Study (LIS) and the World Bank's Povcal data (both as available in March 2021), two widely used sources for survey data. The LIS is generally more focused on advanced economies whereas Povcal also collects survey data for a large number of developing economies. A third inequality data source is the World Inequality Database (WID). Gini values in the WID are computed with information from both survey and fiscal data.

For completeness, we also compare our estimates to the Gini data from the World Income Inequality Database (WIID, Version May 2020) with the caveat that these observations are part of our first stage. Whereas an earlier version of the LIS (from 2019) is among the selected sources from the WIID, we make no use of any version of Povcal or WID data in our first stage estimation.²⁰

²⁰Many more inequality data sets are available such as Solt's Standardized World Income Inequality

Table VI reports the correlation coefficient of our estimates with the inequality values in each of the four established data sets, along with summary statistics of the (absolute) deviation of our estimates from the values found in the respective data sets. These statistics are shown for both levels and period-on-period changes of the Gini. The results resemble those of the 10-fold cross-validation exercise: We find overall a high, though not perfect correlation.

Overall, our methodology points to somewhat higher inequality relative to the survey-based comparison data series. Where a point of comparison is available, about two-thirds of our estimates lie above the highest value (299 out of 457), a quarter (115/457) lie below the lowest value, while the remaining one tenth (43/457) fall in the range given by the comparison data series. Relative to the WID-Ginis that are computed combining surveys with tax data, our estimates are mostly lower (182/265).

Table VI: Comparison with other inequality data sets

		Correlation with estimate	(Abs.) Deviation estimate-data		
			Mean	Std. dev.	Median
Levels	LIS 2021	0.92	4.9	3.1	5.0
	WDI/Povcal 2021	0.92	3.8	2.5	3.4
	WID 2021	0.68	3.8	3.0	3.2
	WIID 2020 (select.)	0.91	4.1	2.7	3.8
Changes	LIS 2021	0.18	1.4	1.3	1.0
	WDI/Povcal 2021	0.13	1.5	1.4	1.1
	WID 2021	0.14	1.6	1.3	1.4
	WIID 2020 (select.)	0.29	1.9	1.7	1.6

Notes: The table presents the correlation of our estimates with four datasets: LIS (as available March 2021), WDI/Povcal (March 2021), WIID Version May 2020, WID (March 2021); along with summary statistics of the absolute deviation between our estimates and values in the alternative data sets.

The numbers in Table VI hide the substantial discrepancies that even the established data series exhibit. Which data series our estimates matches best varies across countries and over time. The six panels in Figure VI illustrate this point for the USA and five European countries, among them Denmark and Sweden. These are two countries known for collecting detailed information on the income distribution so we expect the surveys and fiscal data to have little measurement error.

For Denmark, our method yields Gini estimates that are very close to the WID-Gini notably until 2008. In the decade thereafter, the two series differ more in terms of levels but trace a similar evolution. In contrast, for Sweden, our estimates lie

Data but, like our estimates, these report imputed inequality data and do not consist of primary (survey) data. In addition, many national statistical agencies carry out their own studies. It is beyond the scope of this paper to give a complete overview of alternative data sources.

above both the survey- and survey-and-tax-based inequality series. While the trend that our estimates describe for Sweden is shared with the survey-based data, the patterns diverge more strongly when fiscal data is taken into account. Our trade-based estimates as well as the tax-and-survey-based values of the WID expose a much higher level of inequality in the US than survey data does. Similar to the case of Denmark (and France, panel (d)), after 2008 the WID and our estimate diverge somewhat. Still, our estimates remain throughout very close in levels to the WID-values.

Our baseline method certainly leaves room for refinements and improvements for some countries as we discuss in the next section. Yet, differences between our Gini-estimates and the survey- or tax-based values need not per se cast doubt on our method. For example, our estimations predict a Gini index of between 55 and 51 for Chile, which is up to 7 points above values reported by Povcal and the WIID (for years 2006-2018). Given the recurring protests in Chile on the topic of (perceived) high levels of inequality, it is not unlikely that our approach captures the Chilean reality better than surveys. A similar point can be made for neighboring Argentina. For Venezuela, our estimates are instead close to the Povcal data at around 50 Gini points throughout the sample period, whereas they deviate markedly from the LIS which records a Gini of 36 in 2011.²¹

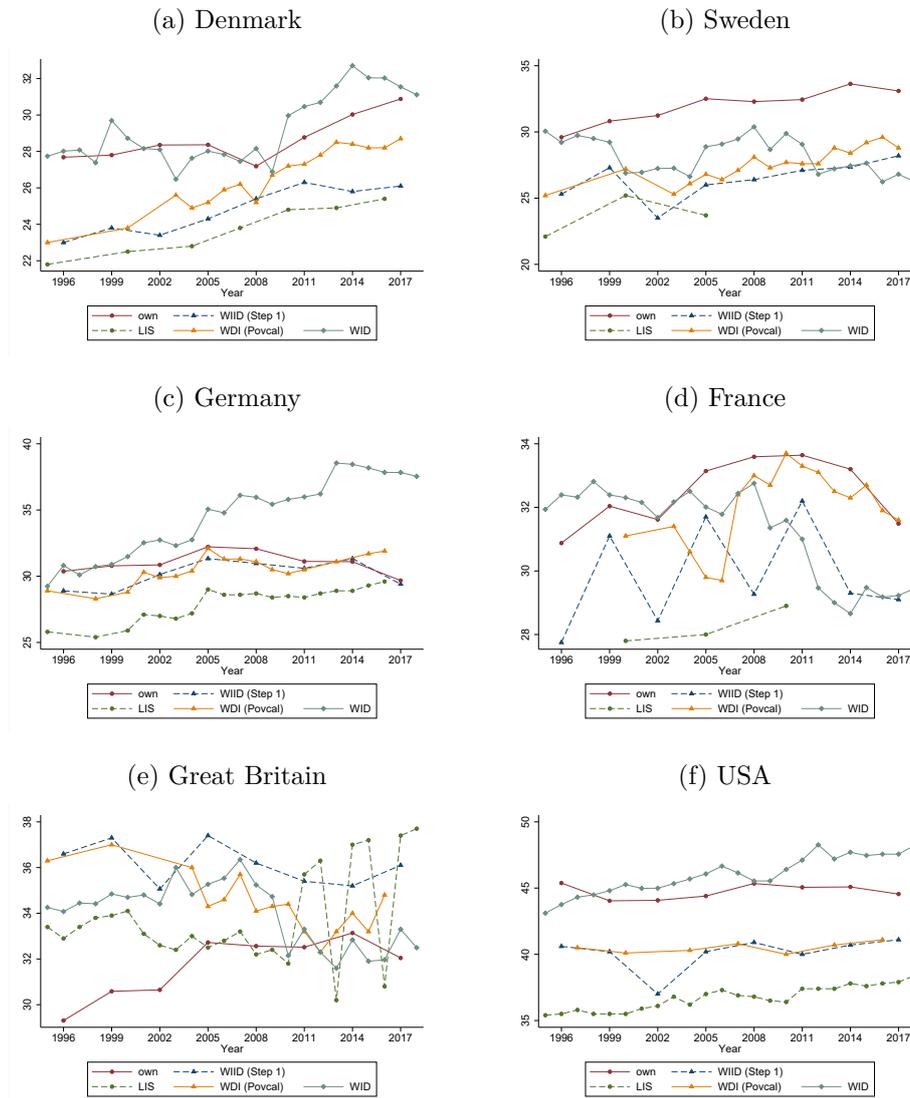
8 Discussion and Robustness

We now discuss how sensitive our results are to a number of modelling assumptions underlying our estimation. Each robustness exercise gives rise to a new vector of β_{jo} and a new set of estimates for inequality. We report in Table VIII the correlation between our estimated Gini coefficients across robustness tests and in Table VII the performance of each exercise compared to the high-quality WIID data as in the previous section.

Time-varying parameters. In our baseline analysis, we allow some preference parameters to vary over time (the spending per good ρ_{jdt} or the idiosyncratic

²¹The LIS has an entry for Venezuela only in the period 2009-2012.

Figure VI: Inequality time series by country - Comparison with four other data sets



Notes: This figures compare the inequality time series per country according to our Gini estimate with the time series from different survey sources. LIS: Luxembourg Income Study; WDI: World Development Indicators; WIID: World Income Inequality Database. The WIID data is used in the first estimation step. Labels on the horizontal axis indicate the middle year of a three-year period.

Table VII: 10-fold cross-validation results by robustness exercise

	First-differences			Level		
	Mean	Sd	Median	Mean	Sd	Median
Time-varying groups	3.73	3.58	2.76	5.36	4.43	4.34
CES nests	2.40	2.17	1.92	4.40	3.30	3.79
Pricing to market	2.63	2.35	2.06	4.86	3.68	4.20
Savings	3.16	2.72	2.53	4.71	3.38	4.02
Long quality ladder products	2.53	2.39	1.73	5.03	3.70	4.73
Poor's consumption basket	2.82	2.39	2.31	5.05	3.83	4.50

Notes: The table presents summary statistics of the distribution of the absolute difference between the out-of-sample predicted Gini and the WIID Gini (in first-differences and levels).

Table VIII: Correlations of in-sample Gini estimates - baseline with alternatives

	Corr. coeff. with baseline
Time-varying parameters	0.72
CES nests	0.99
Pricing to market	0.97
Long quality ladder	0.90
Poor's consumption	0.90
Savings*	0.70

Notes: Correlation coefficients by robustness exercises. *We report the correlation coefficient of our baseline estimates with those of the "savings"-version for completeness despite the conceptual difference between consumption and income inequality.

taste parameter α_{jodt}) while we assume that our main parameters of interest, β_{jo} , are constant over time. These parameters could however have changed over the period. We replicate our analysis by splitting the estimation in two sub-periods, allowing all coefficients to vary by decade. We also allow groups of countries g_j , which were based on the average expenditure vector over the entire sample period in the baseline, to vary by subperiods.

The Gini estimates from this exercise are positively correlated with our baseline estimates. The out-of-sample performance is somewhat weaker than in our baseline, notably concerning the period-on-period changes in inequality. More Gini estimates are below 20, which likely is due to smaller number of observations on which each parameter is identified when the sample is split.

The CES structure. CES preferences imply that the elasticity of substitution is the same between all varieties of a good. The ratio of consumption of any two varieties of a good is thus independent of the consumption of other varieties which implies, inter alia, that not observing the consumption of domestic varieties does not affect our results. We relax this assumption by allowing for a nested CES utility with two nests per product. The elasticity of substitution between any two varieties of the same nest is then equal, but the elasticity across nests is allowed to differ from that within nests. For example, such a structure allows for a different substitution elasticity between high-quality varieties than between high- and low-quality varieties. In practice, it requires that we replace the product-destination-time fixed effect by a nest-product-destination-time fixed effect. We replicate our analysis using varieties from OECD and non-OECD countries as two separate nests.

The second entry in the first column of Table VIII shows that this additional flexibility in our model does not affect our results much. The Gini estimates are highly correlated with the ones from our baseline model and the out-of-sample performance is similar.

Pricing to market. A higher income level in a destination may be associated with higher consumption prices if firms price to market (Simonovska, 2015) or if distribution costs depend on income. Even within countries, firms may charge higher prices to richer consumers and the distribution costs may be correlated with income if, for example, poorer households live in more remote places. While our model abstracts from such concerns by assumption (perfect competition and cost structure), such correlations could bias our estimates by making the error term dependent on income. If all producers of a good j charge a similarly higher price in richer destinations, this is captured by our product-destination fixed effect. There may however be a variety-specific dependency of markups on income, potentially capturing that high- β varieties charge a relatively higher markup in richer countries or to richer consumers within countries.

As we show formally in Appendix B, some of this income-related variation in prices can be accommodated by our framework. To the extent that income affects prices and thus consumption in a systematic way, it is part of the variation of interest. Our estimate of β_{jo} would capture the impact of income on consumption that goes both through preferences and prices. To limit the role of pricing to market, we recompute inequality based on the 182 products (half of the initial number) which have on average the lowest dispersion in unit values within varieties²².

We find once again a high correlation among the Gini estimates based on the thus restricted set of products and the baseline estimates. Additionally, the out-of-sample performance does not deviate much from the baseline results.

Definition of varieties. Our analysis relies on a definition of varieties in the Armington sense, e.g. German cars are a variety. However, in practice, German

²²We compute the variance of log unit values across destinations for each variety and select products with the lowest average variance across varieties, where the average is weighted by the world exports of the variety.

cars consist of different brands, with potentially different preference shifters or income elasticities. German car exports to richer consumers may thus consist mostly of Porsche, and of Opel to poorer consumers. Since the import mix will likely be correlated with the destination's income, this may introduce non-trivial measurement error in our empirical analysis. We show in the appendix B that this would amount to making the error term depend on the quality differentiation within an exporter-product pair and the income of the destination, potentially biasing our estimate. The same restriction on the set of products that we apply to account for pricing to market should partially address the issue of quality differentiation as well: it restricts to products with little variation in unit values across destinations within Armington varieties.

Savings. Our model assumes that households spend their whole income on consumption. In practice, some households save a fraction of their income and the propensity to save may be correlated with income. In the appendix, we show that our strategy remains valid with different savings rates across the income distribution as long as the elasticity of consumer spending to income is constant. If this is not so, our method yields a measure that is conceptually closer to consumption inequality. We therefore replicate our whole analysis using an inequality measure based on consumption rather than income, and show in Table V that it performs equally well.

Using only imports. In our model, all households consume all varieties. Under this assumption, observing imports gives us information about individuals in the whole income distribution. One may be concerned that subsistence farmers in sub-Saharan Africa may not buy Swiss watches but our model predicts that spending on varieties with a high income elasticity approaches zero (even if it is not zero) among the very poor. This could be an issue for our strategy if the very poor do not consume imported varieties of most products as then imports do not convey any information about the poor's income. On the other side of the income distribution, if the very rich consume some goods in their foreign residence, this

will not be reflected in the import patterns of the country.²³ To address these concerns, we run our analysis using a set of goods that are most likely consumed and imported by households throughout the income distribution.

The first set of goods are those with a “long quality ladder” according to Khandelwal (2010) and where “long” means to have a quality ladder length above the median.²⁴ Goods with long quality ladders are those for which very high and very low quality is exported worldwide. Thus, these goods are potentially affordable to and consumed by very diverse groups of consumers.

As an alternative approach we build on Banerjee and Duflo (2007), who summarize insights from surveys about spending of households with less than 1\$, respectively less than 2\$, per day. We select those HS 4 codes that correspond to products mentioned in Banerjee and Duflo (2007). Among these products are food items, alcoholic beverages, tobacco, medication, personal care, textiles, jewelry and watches, radio receivers and TVs, and bicycles.²⁵ These goods are likely also consumed by households across the entire income distribution and therefore may carry the most relevant information for our purposes.

Refining the set of products for the second stage does not have strong effects on the average out-of-sample performance of our method. The Gini estimates are highly correlated to those that are estimated using the full set of final consumption goods. Similarly, the key statistics summarizing the out-of-sample performance are close to the baseline results.

A closer look at the in-sample results however exposes heterogeneous effects across the seven world regions. Figure VII plots once again the evolution of the population-weighted average inequality by world region similar to Figure V. While there are little to no effects on the results for European/Central Asian and Latin American countries, the results for the MENA countries appear highly sensitive to restricting the set of goods from which to estimate inequality. The “quality

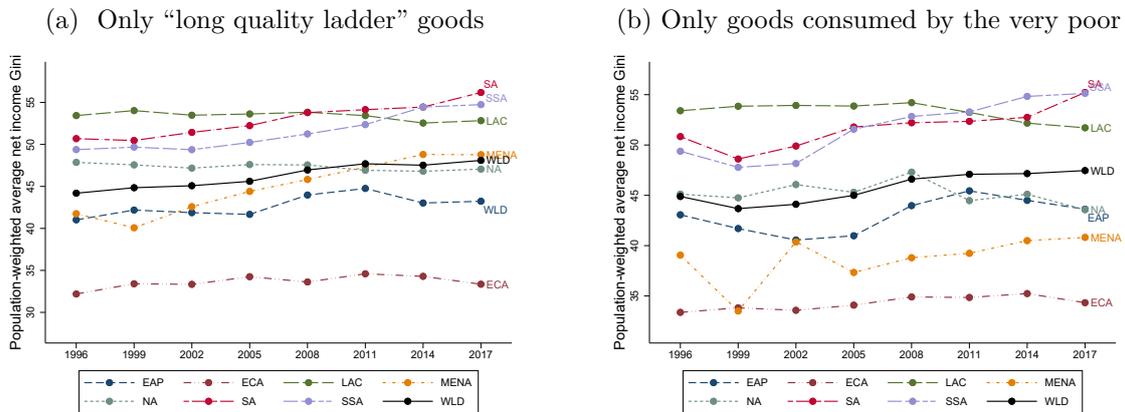
²³Note that if rich individuals buy luxury goods abroad and bring them back home, they will enter customs data too, except of course in case of smuggling. Note also that we use aggregated trade data over 3 years, meaning that the concern is that, over 3 years, the very poor or very rich never spend on imported goods.

²⁴Khandelwal (2010) computes quality as a residual from regressing market shares on prices and a set of fixed effects.

²⁵The list of HS 4 codes is provided in Table XII.

ladder”-restriction leads to estimates that show an even more pronounced increase in average inequality in the MENA region. Instead, restricting to products listed in Banerjee and Duflo (2007) moves inequality down and results in a different path of inequality with a decline of inequality from 1999 to 2005. This second restriction puts however relatively much weight on tobacco and alcoholic beverages and so may not be appropriate for countries with a large Muslim population. We leave more nuanced refinements for future research.

Figure VII: Inequality time series by world region - Refinements



Notes: Population-weighted average Gini estimate by regions: EAP - East Asia & Pacific, ECA - Europe & Central Asia, LAC - Latin America & Caribbean, MENA - Middle East & North Africa, NA - North America, SA - South Asia, SSA - Sub-Saharan Africa, WLD - World average. Labels on the horizontal axis indicate the middle year of a three-year period.

9 Conclusion

In this paper, we propose a novel method for measuring inequality in a country based on its imports of consumer goods. We rely on a model that embeds non-homothetic preferences in an otherwise standard model of international trade à la Armington, in which goods are differentiated by country of origin. This setup generates a non-homothetic gravity equation, where the imports of a variety by a destination country depend on its average income and on the Gini coefficient of its income distribution.

Our empirical approach follows two steps. In the first step, we use a subsample of destinations with high-quality data on inequality and estimate for each variety

the extent to which it is imported by richer or more unequal countries. In the second step, we reverse the logic and estimate, for all countries, the Gini coefficient that best fits the observed import patterns given the variety-specific income elasticities estimated in the first step.

The key advantage of our method is that it relies on international trade data, which are publicly and consistently recorded for virtually all countries. We apply the same method to map trade into inequality for all countries, substantially expanding the coverage of traditional datasets among developing countries. Different cross-validation exercises suggest that our method compares well to existing high-quality data on inequality. Our approach relies on a set of assumptions, which play a key role in mapping import patterns to inequality. We experiment with a number of robustness checks, modifying some of those assumptions, and show that our results are very stable.

References

- Aguiar, M. and M. Bilal (2015). Has consumption inequality mirrored income inequality? *The American Economic Review* 105(9), 2725–2756.
- Almås, I. (2012). International income inequality: Measuring PPP bias by estimating Engel curves for food. *American Economic Review* 102(2), 1093–1117.
- Banerjee, A. V. and E. Duflo (2007). The economic lives of the poor. *Journal of economic perspectives* 21(1), 141–168.
- Besley, T. and T. Persson (2014). Why do developing countries tax so little? *Journal of Economic Perspectives* 28(4), 99–120.
- Blumenstock, J., G. Cadamuro, and R. On (2015). Predicting poverty and wealth from mobile phone metadata. *Science* 350(6264), 1073–1076.
- Caron, J., T. Fally, and J. R. Markusen (2014). International trade puzzles: a solution linking production and preferences. *The Quarterly Journal of Economics* 129(3), 1501–1552.
- Choi, Y. C., D. Hummels, and C. Xiang (2009). Explaining import quality: The role of the income distribution. *Journal of International Economics* 77, 265–275.
- Comin, D. A., D. Lashkari, and M. Mestieri (2015). Structural change with long-run income and price effects. Technical report, National Bureau of Economic Research.
- Deininger, K. and L. Squire (1996). A new data set measuring income inequality. *The World Bank Economic Review* 10(3), 565–591.
- Egger, P. and M. Larch (2008). Interdependent preferential trade agreement memberships: An empirical analysis. *Journal of International Economics* 76(2), 384–399.
- Faber, B. and T. Fally (2017). Firm heterogeneity in consumption baskets: Evidence from home and store scanner data. Technical report, National Bureau of Economic Research.

- Fajgelbaum, P. D., G. M. Grossman, and E. Helpman (2011). Income distribution, product quality, and international trade. *Journal of Political Economy* 119(4), 721–765.
- Fajgelbaum, P. D. and A. K. Khandelwal (2016). Measuring unequal gains from trade. *Quarterly Journal of Economics* 131(3), 1113–1180.
- Feenstra, R. C. and J. Romalis (2014). International prices and endogenous quality. *Quarterly Journal of Economics* 129(2), 477–527.
- Ferreira, F. H., N. Lustig, and D. Teles (2015). *Appraising cross-national income inequality databases: An introduction*. The World Bank.
- Fieler, A. C. (2011). Nonhomotheticity and bilateral trade: Evidence and a quantitative explanation. *Econometrica* 79(4), 1069–1101.
- Galbraith, J. K., B. Halbach, A. Malinowska, A. Shams, and W. Zhang (2014). Utip global inequality data sets 1963-2008: updates, revisions and quality checks.
- Galbraith, J. K. and H. Kum (2005). Estimating the inequality of household incomes: a statistical approach to the creation of a dense and consistent global data set. *Review of Income and Wealth* 51(1), 115–143.
- Gaulier, G. and S. Zignago (2010). Baci: international trade database at the product-level (the 1994-2007 version).
- Hallak, J. C. (2006). Product quality and the direction of trade. *Journal of International Economics* 68, 238–265.
- Hallak, J. C. (2010). A product-quality view of the linder hypothesis. *The Review of Economics and Statistics* 92(3), 453–466.
- Handbury, J. (2019). Are poor cities cheap for everyone? non-homotheticity and the cost of living across U.S. cities. mimeo.
- Hillrichs, D. and G. Vannoorenberghe (2021). Trade costs, home bias and the unequal gains from trade. LIDAM Discussion Paper.

- Ianchovichina, E., L. Mottaghi, and S. Devarajan (2015). Inequality, uprisings, and conflict in the arab world. *MENA Economic Monitor*.
- Khandelwal, A. (2010). The long and short (of) quality ladders. *Review of Economic Studies* 77, 1450–1476.
- Kongsamut, P., S. Rebelo, and D. Xie (2001). Beyond balanced growth. *The Review of Economic Studies* 68(4), 869–882.
- Lessmann, C. and A. Seidel (2017). Regional inequality, convergence, and its determinants—a view from outer space. *European Economic Review* 92, 110–132.
- Linder, S. B. (1961). *An Essay on Trade and Transformation*. Ph. D. thesis, Uppsala.
- Matsuyama, K. (2019). Engel’s law in the global economy: Demand-induced patterns of structural change, innovation, and trade. *Econometrica* 87(2), 497–528.
- McGregor, T., B. Smith, and S. Wills (2019, 07). Measuring inequality. *Oxford Review of Economic Policy* 35(3), 368–395.
- Simonovska, I. (2015). Income differences and prices of tradables: Insights from an online retailer. *The Review of Economic Studies* 82(4), 1612–1656.
- Solt, F. (2009). Standardizing the world income inequality database. *Social Science Quarterly* 90(2), 231–242.

A Derivation of the estimating equation

A.1 Derivation of equation (2)

We use a second-order Taylor approximation of $x_{jo}(\mathcal{B}, \mathbf{A}_d, I_h)$ around $I_h = I_d$:

$$\begin{aligned} x_{jo}(\mathcal{B}, \mathbf{A}_d, I_h) &\approx x_{jo}(\mathcal{B}, \mathbf{A}_d, I_d) + (I_h - I_d) \left. \frac{\partial x_{jo}(\mathcal{B}, \mathbf{A}_d, I_h)}{\partial I_h} \right|_{I_h=I_d} \\ &\quad + \frac{1}{2}(I_h - I_d)^2 \left. \frac{\partial^2 x_{jo}(\mathcal{B}, \mathbf{A}_d, I_h)}{\partial^2 I_h} \right|_{I_h=I_d} \end{aligned} \quad (17)$$

Summing up over all individuals gives:

$$X_{jod} \approx N_d x_{jo}(\mathcal{B}, \mathbf{A}_d, I_d) \left[1 + \frac{1}{2} \left. \frac{\partial^2 x_{jo}(\mathcal{B}, \mathbf{A}_d, I_h)}{\partial^2 I_h} \right|_{I_h=I_d} \frac{I_d^2}{x_{jo}(\mathcal{B}, \mathbf{A}_d, I_d)} G_d \right] \quad (18)$$

where $G_d \equiv \left(\sum_h \frac{N_{dh}}{N_d} (I_h - I_d)^2 \right) / I_{dt}^2$ is the squared coefficient of variation of income in country d . Taking logs and using the approximation $\text{Log}(1+y) \approx y$ on the square bracket gives (2).

A.2 Utility maximization

Maximizing utility with respect to c_{jo} gives the first-order condition:

$$U_d \frac{\rho_{jd}}{C_j} \varphi_{jod}(C_0)^{\frac{1}{1-\gamma_j}} c_{jo}^{\gamma_j-1} C_j^{1-\gamma_j} - \lambda p_{jod} = 0 \quad (19)$$

Taking to the power γ_j , multiplying by $\varphi_{jod}(C_0)$, summing over o and taking to the power $\frac{1}{\gamma_j}$ gives:

$$P_j C_j = \rho_{jd} \frac{U_d}{\lambda} \quad (20)$$

where:

$$P_j = \left[\sum_o \varphi_{jod}(C_0) p_{jod}^{\gamma_j} \right]^{\frac{1}{\gamma_j}}. \quad (21)$$

Plugging back in (19), taking to the power $\gamma_j - 1$ and rearranging gives the optimal spending on variety jo in country d by an individual with consumption C_0 of the

homogeneous good:

$$x_{jod}(C_0) = \frac{\varphi_{jod}(C_0)p_{jod}^{\gamma_j}}{\sum_{o'} \varphi_{j o' d}(C_0)p_{j o' d}^{\gamma_j}} P_j C_j. \quad (22)$$

The first order condition with respect to C_0 is:

$$\frac{U_d}{C_0} \left[\rho_{0d} - \sum_{j=1}^N \frac{\rho_{jd}}{\gamma_j} C_j^{\frac{\gamma_j}{\gamma_j-1}} \sum_o \varphi_{jod}(C_0)^{\frac{1}{1-\gamma_j}} c_{jo}^{\frac{\gamma_j}{\gamma_j-1}} \phi_{jod}(C_0) \right] = \lambda \quad (23)$$

where $\phi_{jod}(C_0) = \frac{\partial \varphi_{jod}(C_0)}{\partial C_0} \frac{C_0}{\varphi_{jod}(C_0)}$. The second part of the square bracket captures the effect of changing the consumption of the homogeneous good on the marginal utility from the consumption of all differentiated varieties. Plugging in the first order condition for c_{jo} gives:

$$\frac{U_d}{C_0} [\rho_{0d} - \Psi_d(C_0)] = \lambda \quad (24)$$

with:

$$\Psi_d(C_0) = \sum_{j=1}^N \frac{\rho_{jd}}{\gamma_j} \frac{\sum_o \varphi_{jod}(C_0) p_{jod}^{\gamma_j} \phi_{jod}(C_0)}{\sum_o \varphi_{jod}(C_0) p_{jod}^{\gamma_j}} \quad (25)$$

$$= \sum_{j=1}^N \frac{\rho_{jd}}{\gamma_j} \left[\sum_o \frac{\alpha_{jod} C_0^{\beta_{jo}} p_{jod}^{\gamma_j}}{\sum_{o'} \alpha_{j o' d} C_0^{\beta_{j o'}} p_{j o' d}^{\gamma_j}} \beta_{jo} - \sum_o \frac{\alpha_{jod} C_0^{\beta_{jo}}}{\sum_{o'} \alpha_{j o' d} C_0^{\beta_{j o'}}} \beta_{jo} \right] \quad (26)$$

where the second line uses (5). Ψ captures the fact that, when increasing C_0 , a consumer will put a higher preference weight on high β varieties and a lower one on low β varieties. If all varieties of j had the same price, the two would exactly cancel out and the choice of C_0 would have no effect on the chosen level of C_j . If prices differ across varieties, however, the individual has an additional incentive to adjust C_0 to give a higher preference weight to varieties with a low price. As long as there is no strong and systematic correlation between $p_{jod}^{\gamma_j}$ and β_{jo} , however, we can assume that $\Psi_{jd}(C_0) = 0$. In the case that high β_{jo} varieties are more costly to produce, making the price p_{jod} proportional to β_{jo}^η , it would suffice to assume that preferences are such that $\varphi_{jod} = \frac{\alpha_{jod} C_0^{\beta_{jo}}}{\sum \alpha_{jod} \beta_{jo}^\eta C_0^{\beta_{j o'}}$ and a zero correlation between the other components of the price and β_{jo} to make $\Psi_d(C_0) = 0$. Combining with

(20), the budget constraint and using that $\rho_{0d} + \sum_{j=1}^N \rho_{jd} = 1$, we obtain, when assuming that $\Psi_d(C_0) = 0$:

$$C_{0d} = \frac{\rho_{0d} - \Psi_d(C_{0d})}{1 - \Psi_d(C_{0d})} I = \rho_{0d} I \quad (27)$$

$$P_{jd} C_{jd} = \frac{\rho_{jd} I}{1 - \Psi_d(C_{0d})} = \rho_{jd} I. \quad (28)$$

As an alternative to imposing that $\Psi_{jd}(C_0) = 0$, we could of course make the assumption that consumers are fully rational and do not take into account that the consumption of the homogeneous good affects their utility from the consumption of differentiated goods.

A.3 Derivation of equation (9)

To simplify the analysis, we denote the income elasticity of x_{jod} as $\lambda_{x_{jod}}$. From (7), we know that:

$$\lambda_{x_{jod}}(I) = 1 + \beta_{jo} - \bar{\beta}_{jd}(I). \quad (29)$$

By definition:

$$\frac{\partial x_{jod}}{\partial I} = \frac{x_{jod}}{I} \lambda_{x_{jod}}. \quad (30)$$

Differentiating with respect to I , rearranging and multiplying by I^2/x_{jod} gives:

$$\frac{\partial^2 x_{jod}}{\partial^2 I} \frac{I^2}{x_{jod}} = (\lambda_{x_{jod}} - 1) \lambda_{x_{jod}} + \frac{\partial \lambda_{x_{jod}}}{\partial I} I. \quad (31)$$

Differentiating $\lambda_{x_{jod}}$ with respect to income gives:

$$\frac{\partial \lambda_{x_{jod}}}{\partial I} I = - \sum_{o'} \frac{\partial s_{j'o'd}}{\partial I} I \beta_{j'o'} = - \sum_{o'} s_{j'o'd}(I) (\lambda_{x_{j'o'd}} - 1) \beta_{j'o'}. \quad (32)$$

Plugging back in (31) gives:

$$\begin{aligned} \frac{\partial^2 x_{jod}}{\partial^2 I} \frac{I^2}{x_{jod}} &= \left(\beta_{jo} - \sum_{o'} s_{j'o'd}(I) \beta_{j'o'} \right) \left(1 + \beta_{jo} - \sum_{o'} s_{j'o'd}(I) \beta_{j'o'} \right) \\ &\quad + \left(\sum_{o'} s_{j'o'd}(I) \beta_{j'o'} \right)^2 - \sum_{o'} s_{j'o'd}(I) \beta_{j'o'}^2. \end{aligned} \quad (33)$$

B Model extensions

We extend the baseline model to incorporate formally some of our robustness checks. We allow for *price discrimination* by letting the price of a variety jo be a function of the income of the consumer (if there is price discrimination within destination or different costs of distribution to individuals with different incomes within a country) and on the average income of the country (if price discrimination happens across countries or if distribution costs are correlated with the average income per capita). We also introduce the possibility of *differentiation within Armington varieties* and use ν to index a “sub-variety”, i.e. a variety of a good within jo ²⁶. We define n_{jo} as the number of sub-varieties within jo . Finally, we allow for *savings* in a reduced form manner and assume that consumers spend a function $f_d(I)$ of their income on goods consumption. Under these assumptions, the exports of product j by country o to destination d become:

$$X_{jod} = \sum_h N_h \sum_{\nu \in jo} x_{jod}(I_h, \beta_\nu), \quad (34)$$

where $x_{jod}(I, \beta_\nu)$ denotes the consumption of a sub-variety of jo with β_ν by an individual with income I and is given by:

$$x_{jod}(I, \beta_\nu) = \frac{\alpha_{jod}(f_d(I))^{\beta_\nu} (p_{jod}(I))^{\gamma_j}}{\sum_o \alpha_{jo'd}(f_d(I))^{\beta_\nu} (p_{jo'd}(I))^{\gamma_j}} \rho_{jd} f_d(I). \quad (35)$$

We assume that $p_{jod}(I) = \tau_{od} c_{jo} m_{jod}(I_d) \mu_{jod}(I)$ to capture that prices can depend on the average income per capita in d and on individual income in line with the discussion on pricing above.

We first perform a second-order Taylor approximation of $x_{jod}(I, \beta_\nu)$ with respect to I around its average value I_d (as in section 3.1) and with respect to β_ν around its average value β_{jo} and sum up over all individuals and sub-varieties within jo to

²⁶We could also make the price of a sub-variety depend on β_ν , for example if higher quality variety are more expensive to produce and preferred by the rich. This would simply add a few terms without affecting the main message of this section.

obtain:

$$\frac{X_{jod}}{N_d n_{jo}} = x_{jod}(I_d, \beta_{jo}) + \frac{1}{2} \frac{\partial^2 x_{jod}(I, \beta_{jo})}{\partial^2 I} \Big|_{I=I_d} I_d^2 G_d + \frac{1}{2} \frac{\partial^2 x_{jod}(I_d, \beta)}{\partial^2 \beta} \Big|_{\beta=\beta_{jo}} V_{jo} \quad (36)$$

where N_d is the number of individuals in destination d , n_{jo} is the number of sub-varieties in jo and V_{jo} is the variance of β_ν within jo .

We denote λ_y as the elasticity of variable y with respect to I , when evaluated at $I = I_d$ and $\beta_\nu = \beta_{jod}$ and show that:

$$\lambda_{x_{jod}} = \lambda_{f_d} \left(1 + \beta_{jo} - \sum_{o'} s_{jo'd} \beta_{jo'} \right) + \gamma_j \left(\lambda_{\mu_{jod}} - \sum_{o'} s_{jo'd} \lambda_{\mu_{jo'd}} \right). \quad (37)$$

where $s_{jod} = n_{jo} x_{jod}(I_d, \beta_{jo}) / (\rho_{jd} I_d)$ is the share of income spent on all sub-varieties from jo if they had β_{jo} by an individual with income I_d . We define:

$$B_{jod} \equiv \beta_{jo} \lambda_{f_d} + \gamma_j \lambda_{\mu_{jod}}, \quad (38)$$

which allows rewriting:

$$\lambda_{x_{jod}} = B_{jod} - \sum_{o'} s_{jo'd} B_{jo'd} + \lambda_{f_d} \quad (39)$$

Differentiating $\lambda_{x_{jod}}$ with respect to I gives:

$$\begin{aligned} \frac{\partial \lambda_{x_{jod}}}{\partial I} I &= (\beta_{jo} + 1 - \sum_{o'} s_{jo'd} \beta_{jo'}) \frac{\partial \lambda_{f_d}}{\partial I} I + \gamma_j \left(\frac{\partial \lambda_{\mu_{jod}}}{\partial I} I - \sum_{o'} s_{jo'd} \frac{\partial \lambda_{\mu_{jo'd}}}{\partial I} I \right) \\ &- \sum_{o'} \frac{\partial s_{jo'd}}{\partial I} I B_{jo'd} \end{aligned} \quad (40)$$

Collecting terms and using (31) shows that:

$$\begin{aligned} \frac{\partial^2 x_{jod}}{\partial^2 I} \frac{I_d^2}{x_{jod}} \Big|_{I=I_d} &= (B_{jod} - \bar{B}_{jd} + \lambda_{f_d})(B_{jod} - \bar{B}_{jd} + \lambda_{f_d} - 1) - \sum_{o'} s_{jo'd} B_{jo'd}^2 \\ &+ \bar{B}_{jd}^2 + (1 - \lambda_{f_d}) \bar{B}_{jd} \\ &+ (\beta_{jo} + 1 - \bar{\beta}_{jd}) \frac{\partial \lambda_{f_d}}{\partial I} I \\ &+ \gamma_j \left(\frac{\partial \lambda_{\mu_{jod}}}{\partial I} I - \sum_{o'} s_{jo'd} \frac{\partial \lambda_{\mu_{jo'd}}}{\partial I} I \right) \end{aligned} \quad (41)$$

where $\bar{\beta}_{jd} \equiv \sum_{o'} s_{j'o'd} \beta_{j'o'}$ and $\bar{B}_{jd} \equiv \sum_{o'} s_{j'o'd} B_{j'o'd}$. The above expression boils down to (33) under our baseline assumptions: $\lambda_{\mu_{jod}} = 0$ and $\lambda_{f_d} = 1$.

Differentiating $x_{jod}(I_d, \beta_\nu)$ twice with respect to β_ν and evaluating it at $\beta_\nu = \beta_{jo}$ shows that:

$$\frac{\partial^2 x_{jod}(I_d, \beta_{jo})}{\partial^2 \beta_{jo}} \frac{1}{x_{jod}(I_d, \beta_{jo})} = \left(1 - \frac{s_{jod}}{n_{jo}}\right) \left(1 - 2 \frac{s_{jod}}{n_{jo}}\right) (\text{Log}(f_d(I_d)))^2 \equiv \Delta_{jod}(I_d) \quad (42)$$

Plugging the previous results in (36), taking logs and using that $\text{Log}(1+y) \approx y$ for small y , we obtain the following gravity equation, where Γ_{jo} and Γ_{jd} collect all the jo and jd specific terms:

$$\begin{aligned} \text{Log}(X_{jod}) &\approx \text{Log}(\alpha_{jod}) + \gamma_j \text{Log}(\tau_{od}) + \Gamma_{jo} + \Gamma_{jd} + \gamma_j \text{Log}(m_{jo}(I_d) \mu_{jo}(I_d)) \\ &+ \beta_{jo} \text{Log}(f_d(I_d)) \\ &+ \frac{1}{2} \left[B_{jod} (2\lambda_{f_d} - 1 + B_{jod} - \bar{B}_{jd}) + \gamma_j \frac{\partial \lambda_{\mu_{jod}}}{\partial I_d} I_d \right] G_d \\ &+ \Delta_{jod}(I_d) V_{jo} \end{aligned} \quad (43)$$

The last term captures the effect of *quality differentiation within Armington varieties*. It is a function of β_{jo} as well as of I_d and may thus bias our β_{jo} coefficient estimates of the first stage. In our robustness check, we thus concentrate on products for which we think that V_{jo} is reasonably small, i.e. where the heterogeneity in β_ν within Armington varieties is small. Note that differentiation within Armington varieties is only an issue if it is associated with different β_ν , i.e. vertical differentiation. Purely horizontal differentiation would be captured by the fixed effect Γ_{jo} and would not bias our estimates. *Price discrimination* enters through different terms in equation (43). To the extent that prices depend only on the average income of a destination, the effects of price discrimination would only enter through the term $m_{jod}(I_d)$. Price discrimination based on income within destinations is captured by all the terms with a μ , including B . It is worth pointing out the case where $m_{jo}(I_d) = 1$ and $\mu_{jo}(I) = I^{\delta_{jo}}$. In that case, and if $f_d(I_d) = I_d$ as in our baseline specification, the equation boils down to our baseline gravity equation replacing β_{jo} by B_{jo} . In that case, we estimate the effect of income on

consumption that runs through both preferences and prices. Our strategy relies on capturing the link between income and consumption, no matter through which exact channel, and pricing to market can thus be part of that variation of interest for our method. To the extent that pricing to market enters in a different functional form and depends on β_{jo} , it does however add a component to our error term that depends both on β_{jo} and on income. In section 8, we thus reestimate our model on a subset of products for which we think that it is less of a concern. Finally, we allow for the *share of income spent on consumption* to vary with income, with total spending being a function $f_d(I)$. If $f_d(I) = \phi_d I^\psi$, for example, this does not affect our strategy and our results apply even when households save some of their income. In more general cases, it may bias our estimate. In that case, our measure may be interpreted as a measure of consumption inequality, which is also available for some countries. In section 8, we thus recompute our analysis when using consumption instead of income inequality as a concept.

C Country Groups

Our methodology rests on the assumption that the preference parameters are constant within HS 4-digit products across importer countries with and without Gini data. We group countries and claim that preferences of countries with and without Gini data within groups are sufficiently similar within HS 4-goods to infer inequality from import patterns.

The grouping procedure is theory-based. Notably, if the income elasticity β_{jo} is constant across households and countries, inequality affects imports of a variety jo differently across importer countries only if the weighted average income elasticity of the countries' consumption baskets $\bar{\beta}_{jd}$ differs. See the last term in equation (11). The weighted average income elasticity is defined in equation (8), which we repeat here for convenience:

$$\bar{\beta}_{jd}(I_h) \equiv \sum_o s_{jod}(I_h)\beta_{jo}. \quad (44)$$

Define the \mathbf{s}_{dj} as the row vector collecting expenditure shares by country d in product j from the origin countries o . Similarly, define β_j as the row vector of income elasticities of product j . Then we can rewrite the weighted average income elasticity in vector notation

$$\bar{\beta}_{jd}(I_h) \equiv \mathbf{s}_{dj}\beta_j'. \quad (45)$$

We group countries with the highest similarity in expenditure share vectors. The goal of our procedure is to form groups of countries with and without Gini data such that those countries included in the first step regression are representative of those that are not included in the regression. "Representative" means that for all countries within a group the income elasticity of the variety jo is either above or below the average such that those countries grouped together share the shape of a product's Engel curve.

We deviate slightly from the theory and choose to calculate expenditure shares at the HS Chapter level (instead of the HS 4-digit level) and over the full sample

period (instead of by period). That is, we calculate the expenditure shares:

$$s_{d,C} = \frac{X_{C,od}}{X_{C,d}} = \sum_{j \in C} \sum_t X_{jodt} / \sum_{j \in C} \sum_t \sum_o X_{jodt}, \quad (46)$$

where C is an HS Chapter and $j \in C$ are all HS 4-digit codes belonging to Chapter C .

We then calculate the cosine similarity for each country pair, i.e. we calculate the cosine of the angle between the expenditure vectors of two countries d and d' . The cosine similarity considers the similarity in the direction of the expenditure vectors as opposed to the length of the vectors (like a Euclidean distance measure would). The formula is

$$sim_{d,d'} = \frac{\mathbf{s}_{d,C} * \mathbf{s}_{d',C}}{\|\mathbf{s}_{d,C}\| \|\mathbf{s}_{d',C}\|} = \frac{\sum_{o=1}^N X_{oCd} X_{oCd'}}{\left(\sum_{o=1}^N X_{oCd}^2 \sum_{o'=1}^N X_{o'Cd'}^2\right)^{1/2}} \quad (47)$$

Next, we use hierarchical clustering using Ward's linkage method - first, for the countries with Gini data in more than half of all sample periods (first step sample). Each country is initially a singular cluster. Countries are then being merged to form groups with minimum within-cluster variance of the similarity index. The process continues until we are left with three groups of countries from the first-step-sample.

We then calculate for each country without Gini data $d \in \mathbb{U}$ the average similarity in expenditure shares with the expenditure shares of the countries with Gini data per group $d' \in g \subset \mathbb{O}$. That is for each country without Gini data, we get three average similarities with countries with Gini data that have already been grouped by Ward's linkage method in the previous step:

$$av. sim_{d,g} = \frac{1}{N} \sum_{d' \in g \subset \mathbb{O}} sim_{d,d'}.$$

We then assign each country without Gini data to that group of countries with which it shares the highest average similarity. The order of forming groups first on the first-step sample and matching the remaining countries later to these groups is important. If we were to form groups with the full sample we would risk that in the extreme all countries with Gini are assigned to one group and all countries

Table IX: Average similarity in expenditure shares within groups

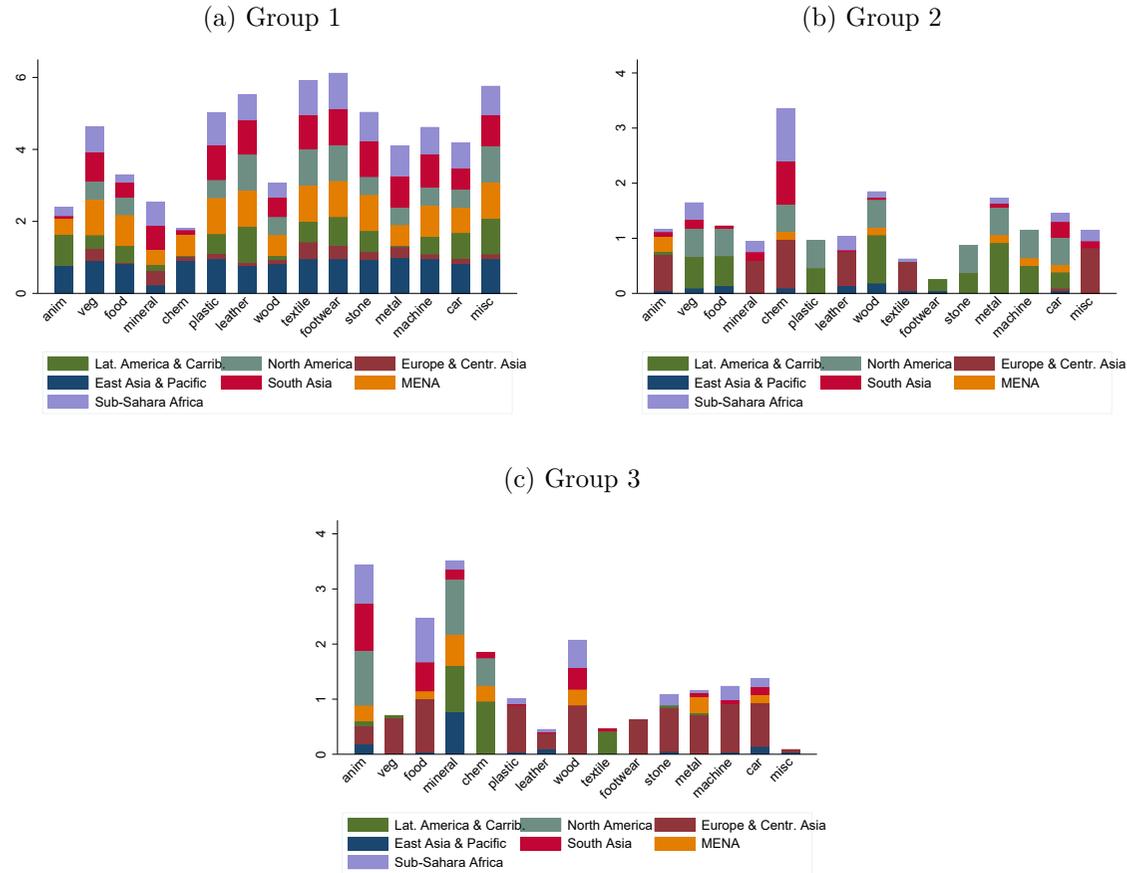
	Group 1	Group 2	Group 3
Mean	0.410	0.354	0.470
Std. dev.	0.204	0.172	0.144
Median	0.382	0.349	0.453
Min	0.068	0.114	0.301
Max	0.704	0.729	0.787
# HS Chapters		15	

Notes: Summary statistics of average similarity within country group across 15 HS Chapters.

without Gini data are assigned to another group.

We provide summary statistics on the within-group similarity in Table IX, and a break-down of the groups by world region in Table X. See also Figure IX for the availability of Gini data by world region.

Figure VIII: Country group composition by HS Chapter



Notes: Share of countries in given group coming from a given region stacked in one bar per HS Chapter.

Table X: Group composition by HS Chapter and World Region

Chapter	Group	EAP	ECA	LAC	MENA	NA	SA	SSA	Total
1	1	14	15	3	16	2	3	36	89
1	2	1	30	1	2	0	2	4	40
1	3	7	0	20	2	0	2	1	32
2	1	20	15	9	15	1	7	34	101
2	2	2	0	14	5	1	0	7	29
2	3	0	30	1	0	0	0	0	31
3	1	18	3	11	3	1	6	16	58
3	2	3	0	13	0	1	0	3	20
3	3	1	42	0	17	0	1	22	83
4	1	3	18	4	12	0	3	28	68
4	2	0	27	0	4	0	0	5	36
4	3	19	0	20	4	2	4	8	57
5	1	18	2	7	0	1	1	1	30
5	2	1	0	17	0	0	2	4	24
5	3	3	43	0	20	1	4	36	107
6	1	21	9	13	18	1	7	40	109
6	2	0	0	11	0	1	0	0	12
6	3	1	36	0	2	0	0	1	40
7	1	16	5	24	13	2	6	34	100
7	2	5	20	0	7	0	1	6	39
7	3	1	20	0	0	0	0	1	22
8	1	20	5	3	9	1	5	23	66
8	2	2	0	21	1	1	1	1	27
8	3	0	40	0	10	0	1	17	68
9	1	21	7	15	18	2	7	38	108
9	2	0	0	9	0	0	0	1	10
9	3	1	38	0	2	0	0	2	43
10	1	21	16	18	20	2	7	41	125
10	2	0	0	4	0	0	0	0	4
10	3	1	29	2	0	0	0	0	32

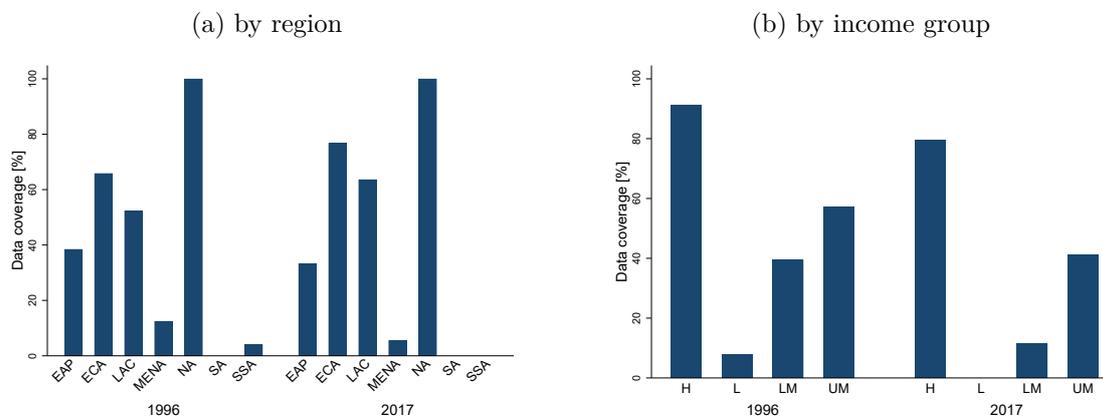
Notes: The table presents the number of countries from a given world region that are selected into the same group per HS Chapter.

Chapter	Group	EAP	ECA	LAC	MENA	NA	SA	SSA	Total
11	1	21	11	14	17	1	7	41	112
11	2	1	0	10	0	1	0	0	12
11	3	0	34	0	3	0	0	0	37
12	1	6	1	23	5	1	1	16	53
12	2	0	29	0	2	0	2	6	39
12	3	16	15	1	13	1	4	19	69
13	1	21	6	11	16	0	7	38	99
13	2	0	0	13	0	2	0	0	15
13	3	1	39	0	4	0	0	3	47
14	1	18	5	18	13	1	5	21	81
14	2	1	2	6	4	1	1	14	29
14	3	3	38	0	3	0	1	6	51
15	1	21	6	24	16	2	7	36	112
15	2	0	37	0	4	0	0	5	46
15	3	1	2	0	0	0	0	0	3

Notes: The table presents the number of countries from a given world region that are selected into the same group per HS Chapter.

D Figures

Figure IX: Availability of net income Gini in WIID



Notes: The chart shows share of countries in a world region/income group for which high quality net income inequality data is available in the World Income Inequality Database. Regions: EAP - East Asia Pacific, ECA - Europe & Central Asia, LAC - Latin America & Caribbean, MENA - Middle East & North Africa, NA - North America, SA - South Asia, SSA - Sub-Sahara Africa. Income groups: H - high income, UM - upper middle income, LM - lower middle income, L - low income.

E Tables

Table XI: Gini availability for welfare definition: Consumption

Years	WIID 2020	trade-based estimate	Povcal 2021
1995-1997	45	144	34
1998-2000	32	148	38
2001-2003	24	151	54
2004-2006	42	150	55
2007-2009	43	155	55
2010-2012	47	154	58
2013-2015	37	154	52
2016-2018	2	153	40
Total	272	1209	386

Notes: World Development Indicators (Povcal) accessed March 2021. WIID: Version May 2020. Gini definition: Consumption inequality.

Table XII: Products consumed by the poor

Product group	HS 4-digit code										
Vegetables	0901	0902	0903	1006	1512	1514					
Foodstuffs	1704	1806	1902	2101	2203	2204	2205	2206	2402	2403	
Medication, personal care	3004	3304	3306	3401							
Plastic tableware	3924										
Cigarette paper	4813										
Textiles	6109	6110	6205	6301	6302	6309					
Jewelery	7018	7113	7114	7117							
Radio/TV	8519	8527	8528								
Bicycles	8712										
Festivity articles, misc.	9404	9505	9608	9609	9613						

Notes: The table lists HS 4 codes corresponding to products mentioned in Banerjee and Duflo (2007) to which we restrict the data for a robustness exercise (section 8).

Table XIII: HS Chapter Description

HS Chapter	Title	Examples
1	Animal & Animal Products	Meat, fish, dairy products
2	Vegetable Products	Vegetables, fruit, plants
3	Foodstuffs	Edible preparations, drinks, tobacco
4	Mineral Products	Salt, chalk
5	Chemicals and Allied Industries	Soap, perfumes, fertilizers
6	Plastics/Rubber	Tableware, kitchenware, apparel of rubber
7	Raw Hides, Skins, Leather & Furs	Apparel and clothing accessoires, saddlery, cases
8	Wood & Wood Products	Paper, paper board, books, printed matter
9	Textiles	Clothes
10	Footwear/Headgear	Shoes
11	Stone/Glass	Pearls, Diamonds, Ceramic Products, Glass
12	Metals	Tools, cutlery
13	Machinery/Electrical	Computer, vacuum cleaner, dish washer
14	Transportation	Cars, bikes
15	Miscellaneous	Sports equipment, music instruments, toys

Table XIV: Trade-based income Gini estimates

ISO-3	1995-1997	1998-2000	2001-2003	2004-2006	2007-2009	2010-2012	2013-2015	2016-2018
AFG [†]					23.5*	42	45.9	49.5
AGO	46.3	49.7	45.7	49.6	49.1	51.2	50.2	49.2
ALB	38	34.8	33	33.7	27.9	28.8	30.4	31.9
ARE* [†]					21.7*	26.7	25.3	27.8
ARG	48.6	50.2	50.5	51.6	51.6	51.7	49.8	48.7
ARM	48.7	37.2	27.8	31.1	30	30.4	33.3	32.7
AUS	42.1	42.9	43	41.5	41.1	41	38.2	39.2
AUT	32.1	33	32	32.3	32.5	33	32.9	31.6
AZE	43.4	39.4	36.6	34.5	33.2	33.6	34.7	31.3
BDI	52.6	48.6	50.6*	47.3	47.3	43.6	47.7	46.9
BEL	28.5	27.7	26.3	27.4	27.2	28.5	28.1	27.1
BEN	46.2*	50.5	51.6	52.2	53.7	50.8	51.7	53.2
BFA	44.1	46.1	45.9	46.4	44.2	43.8	45.8	45.1
BGD	52.4	51.1	51.2	52.5	51.6	52.1	52.8	54.2
BGR	32.1	33.2	32.7	32.1	34.4	39	39.9	40
BHR*			10.7*		15.9*	14.4*	20.7*	26.2
BIH	38.5	37.3	35.9	33.6	29.9	31.1	28.9	29.1
BLR	38.7	32.6	32.5	31.1	30.8	31.2	32.8	29.1
BOL	53.4	55.7	54.6	53.9	53.6	54	53.1	51.3
BRA	58.3	58.5	56.5	55	55.5	55.5	55.4	55.2

Notes: Table presents estimates of net income Gini index. Estimates missing if: [◇] trade data missing, [†] GNI per capita data missing, * estimated coefficient of variation squared negative. Values marked with * if deemed too sensitive to specific products imported.

Table XIV: Trade-based income Gini estimates

ISO-3	1995-1997	1998-2000	2001-2003	2004-2006	2007-2009	2010-2012	2013-2015	2016-2018
BTN	39.5*	8.9*	42.2*	38.5*	37.3*	22.9*	29.5*	26.1*
CAF	43.7*	40.3*	39.3*	39*	37.6*	38.2*	44*	47.5
CAN	36.8	38.2	39.1	39.7	39.9	39.6	38.1	39.9
CHE	27	28.4	29.2	31.3	31.6	30.7	26.8	22.9
CHL	53.2	54.5	54.6	55	53.7	53.1	51.7	51.1
CHN	43	41.4	41.3	41.4	44	44.9	42	41.5
CIV	50.2	51.1	52.6	51.5	53.8	51.6	54.8	51.1
CMR	46.5	45.2	47.8	49	48.7	50.8	52.1	49.6
COD	48.4	43.3	47.4	45.9	47.7	46.1	47.4	48.5
COG	40.8*	46.4	47.2	48.4	47.4	49.6	47.1	51.2
COL	54.8	55.9	55.4	55.6	56.5	55.5	56	54.8
COM	41.6*	47.8	43.2	44.8	44.7	50.8	49.7	50
CPV	45.5	41.7	43.1	41.8	40.9	40.3	39.7	38
CRI	49.5	50.2	50.6	50.6	51.3	50.4	51.6	51.1
CYP	30.6	30.2	29	34.2	35.9	37.6	38.7	37.7
CZE	27.8	27.4	27.3	29.8	31	31.5	31.1	30
DEU	30.4	30.8	30.9	32.2	32.1	31.1	31.1	29.7
DJI [†]							50.5	52.2
DNK	27.7	27.8	28.4	28.4	27.2	28.8	30	30.9
DOM	50.5	52.7	51.2	51.3	51.6	50.3	49.2	49

Notes: Table presents estimates of net income Gini index. Estimates missing if: [◇] trade data missing, [†] GNI per capita data missing, * estimated coefficient of variation squared negative. Values marked with * if deemed too sensitive to specific products imported.

Table XIV: Trade-based income Gini estimates

ISO-3	1995-1997	1998-2000	2001-2003	2004-2006	2007-2009	2010-2012	2013-2015	2016-2018
DZA	42.3	43.8	45.2	45.6	44.5	43.4	44.7	48
ECU	52.3	53.3	55	54.7	54.6	54.3	53.2	54.8
EGY	40.2	40.4	41.3	41.2	41.3	42.9	46.7	44.6
ERI [†]	36.4*	37.7*	38.4	38.6*	36.4*	38.7*		
ESP	33.7	35.5	37	37.1	38.4	38.6	39.1	38.6
EST [†]		33.6	32.2	35.2	35.6	35.3	34.8	34.5
ETH	49	49.8	52.9	53.2	52.7	49.5	51.8	53.1
FIN	28.4	29.6	30.7	32	29.8	32.5	32.4	31.5
FJI	47.3	45.7*	35.4*	44.7	41.8*	39.4*	44.2	47.1
FRA	30.9	32	31.6	33.1	33.6	33.6	33.2	31.5
GAB	34.7	34.5	34.7	36.2	37.5	38.2	36.1	35.1
GBR	29.3	30.6	30.6	32.7	32.6	32.5	33.1	32
GEO	42.6	36.8	35.8	36.7	36.6	39.8	38.8	38.3
GHA	50	51	49.3	53.1	53.4	54	52.6	55.1
GIN	50.6	50.8	52	50.7	51.2	52.6	54.2	53.8
GMB	37*	37.8*	39.5	41.7	39.1	44.6	47.3	48.7
GNB	52.6*	52*	44.2*	51.1*	56.4	56.9	56	59.3
GNQ	29.8*	42.6	39.2	31.5*	24.2*	29	34.2	34.1
GRC [†]				37.4	37.4	37.9	38.5	38.8
GTM	51.2	55.4	51.7	51.8	51.1	51.9	51.7	51.2

Notes: Table presents estimates of net income Gini index. Estimates missing if: [◇] trade data missing, [†] GNI per capita data missing, * estimated coefficient of variation squared negative. Values marked with * if deemed too sensitive to specific products imported.

Table XIV: Trade-based income Gini estimates

ISO-3	1995-1997	1998-2000	2001-2003	2004-2006	2007-2009	2010-2012	2013-2015	2016-2018
GUY	30.1*	32.2*	36.4	36.7	41.6	41.2	40.4	42.5
HKG	34.1	36.2	34.7	33.9	30.8	34.5	33.4	37.4
HND	53.7	54.9	54.5	54.7	54.1	53.7	53.9	53.8
HRV	35.3	34.6	30.6	29.9	30.3	30.6	38.4	39.1
HTI	43.5	47.5	46.4	47.1	45.6	43.8	43.1	46.7
HUN	31.9	31.9	30.8	36.7	37.3	37.7	37.5	36.5
IDN	40.7	43.2	43.8	44.8	46.1	46.3	46.7	48.6
IND	48.5	49.2	50.4	51.8	52.5	52.8	52.9	54.2
IRL	29	28	30.1	31.3	32	31.7	32	31
IRN	29.7*	35	34.9	35.4	38.1	38.5	42.9	44.7
IRQ	45.7*	43.1*	30.2*	32.4*	42.6	42.4	43.5	45.8
ISR	36.1	34.3	35.4	35.9	35	36.3	36.2	33.7
ITA	33.2	34.4	35.5	36.3	36.9	37.4	38.4	37.4
JAM	43.6	42.1	41.7	41.4	46.8	46.6	45.2	45.8
JOR	41.7	37.6	39.6	41.6	42.4	43.3	44.7	46.1
JPN	35.8	41	40.1	40.7	42	40.6	43.2	42.9
KAZ	23.1*	24*	20.4*	15.8*	21.4*	16.9*	14.8*	14.3*
KEN	48	48.4	48.3	48.9	50.6	51	51	51.2
KGZ	41.3	39	37.7	33.9	37.8	38.2	34.4	33
KHM	47.2*	43.7*	43.6*	41.5*	46.1	46.9	48	51.5

Notes: Table presents estimates of net income Gini index. Estimates missing if: \diamond trade data missing, \dagger GNI per capita data missing, * estimated coefficient of variation squared negative. Values marked with * if deemed too sensitive to specific products imported.

Table XIV: Trade-based income Gini estimates

ISO-3	1995-1997	1998-2000	2001-2003	2004-2006	2007-2009	2010-2012	2013-2015	2016-2018
KOR	34.7	35.2	34.6	37.8	38	41.1	39.2	40.4
KWT*							13.5*	23.1
LAO*	44.5*	38.1*	24.9*		34.9*	36.7*	34.9*	35.7*
LBN	29.6	27.6	31.4	31.3	28.8	29.4	29.8	31.6
LBR [†]		42.3	44.5	45	46.5	47.9	49.4	49
LBY [†]			32.6	36.5	34.7	36.2	37.6	38.5
LKA	47.7	47.6	43.2	46.2	47.7	49.3	47.4	49.2
LTU	38	35.7	34.7	37.7	38.4	38.3	38.3	38.4
LVA	37	35.5	35	37.2	37.7	37	37.4	36.5
MAC*								
MAR	41.6	41.3	43.8	44.4	45.2	46.4	47.1	45.6
MDA	42.6	40.3	39	36.6	33.9	33	30.2	30.1
MDG	47	49.2	48.1	45.8	49.5	49.5	52.4	52.2
MEX	49.3	51	52.1	52.4	52	52.2	50.7	49.6
MKD	38.6	36.6	34.4	32	30.6	32.8	31.9	30.3
MLI	48.8	51.7	48.6	51.4	51.3	52.5	54.3	52.6
MMR	49.8	52.9	46.3	46.9	39.3*	48.6	46	50.1
MNG	23.9*	30.5*	32.6*	23.3*	29.6*	21.2*	21.1*	26.5*
MOZ	55	55.1	56	57	55.6	59.7	60.6	59.6
MRT	51.4*	42.2*	43.3*	50.8	48.2*	49.5	53.3	56

Notes: Table presents estimates of net income Gini index. Estimates missing if: [◇] trade data missing, [†] GNI per capita data missing, * estimated coefficient of variation squared negative. Values marked with * if deemed too sensitive to specific products imported.

Table XIV: Trade-based income Gini estimates

ISO-3	1995-1997	1998-2000	2001-2003	2004-2006	2007-2009	2010-2012	2013-2015	2016-2018
MUS	40	37.7	36.9	37.6	39.6	38.9	38.1	36
MWI	47.1	40.5*	40.2*	42.2	45.5	44.9	46.3	46.4
MYS	41.1	43.6	42.3	42	43.7	45.6	45	45.7
NER	41.7*	43.8	49.1	48.5	45.8	46	49.7	52.4
NGA	53.3	53.8	54.7	55.5	54.4	56.6	56.5	57.2
NIC	45.2	50.3	50.6	51.8	51.3	51.1	51.5	50.7
NLD	26.4	25.3	26.3	28.1	25.7	24.6	26.3	24.8
NOR	24.3	25.2	24.9	25.2	23.3	21.4	22.9	22.4
NPL	40.4*	42.4	43.8	42.6	42.9	43.8	43.5	44.7
NZL	40.3	38.1	37.8	41	38.1	38	36	36.4
OMN*	13.6*	14.6*	13.8*		16.2*	24.6*	32	33.5
PAK	51.5	49.9	49.7	54.2	53.5	54.2	54.9	55.2
PAN	56.8	57.6	57.5	57	56.1	56.1	55.3	54.8
PER	54.6	54.2	55.6	55.7	55.6	53.5	54.6	55
PHL	43.7	44.7	45.7	46.1	48.4	47.5	47.4	48.5
PNG	45.5*	45.3*	43*	39.7*	41.3*	50.1	43.4*	45.4*
POL	30.2	31.1	31.3	35.2	36.2	36.8	37.8	36.3
PRT	39.4	40.6	41.6	42	42.4	41.6	41.7	41.2
PRY	52.4	54.2	54.4	54.4	55.5	54.8	53.8	54
PSE		53.3*	45.4*	35*	43.9	45	43	41.5

Notes: Table presents estimates of net income Gini index. Estimates missing if: \diamond trade data missing, \dagger GNI per capita data missing, * estimated coefficient of variation squared negative. Values marked with * if deemed too sensitive to specific products imported.

Table XIV: Trade-based income Gini estimates

ISO-3	1995-1997	1998-2000	2001-2003	2004-2006	2007-2009	2010-2012	2013-2015	2016-2018
QAT [†]		15.8*	16.8*	18.4	16	17	15.6	18.5
ROU	34.6	35.3	35.8	36	39.7	40.8	41.8	40.9
RUS	37.3	37.9	36.9	37.2	35.2	34	33.9	32.9
RWA	41.1*	48.6	45.4	41.5*	46.1	47	49.2	49.2
SAU*		15.9*	14.5*	18*	26.2*	27.3*	30	33.3
SDN [◇]						43.6	48.2	49.7
SEN	48.9	49.9	49.1	49.8	51.3	51.7	52.3	52.1
SGP	28.6*	22.6*	24.6	27	26.9	26.3	26.1	31.4
SLB	45.2*	49.8*	56.3*	36.7*	44.6*	50.2*	45.8*	47.3*
SLE	44*	46.3*	43.7	46.2	50.2	48.4	51.4	54.6
SLV	51.7	52.8	52.4	53.5	51.7	51.6	50.9	50.6
SRB				32.4	33.6	32.9	34	33.6
SUR	25.7*	26.3*	20.6*	26.2*	35.6	34	33	31.6
SVK	27.9	27.8	27.2	30.6	31.1	33.2	33.4	33.5
SVN	30	29.7	29.3	30.6	31	32.2	32.8	31.3
SWE	29.6	30.8	31.2	32.5	32.3	32.4	33.6	33.1
TCD	47.2	33.3*	35.1*	46.8	41.4*	42.9	43.7	44.4
TGO	52	51.2	51.5	56.2	55	56.4	57.5	57.1
THA	36.6	41.7	44.3	42.6	41.8	43.8	44.6	45.4
TJK	51.3	48.2*	40.3*	41.2*	40.1*	38.7*	43.7	43.1*

Notes: Table presents estimates of net income Gini index. Estimates missing if: [◇] trade data missing, [†] GNI per capita data missing, * estimated coefficient of variation squared negative. Values marked with * if deemed too sensitive to specific products imported.

Table XIV: Trade-based income Gini estimates

ISO-3	1995-1997	1998-2000	2001-2003	2004-2006	2007-2009	2010-2012	2013-2015	2016-2018
TKM	42.4	41.4	37.8*	41.1	39.3	39.3	38.1	36.3
TLS [†]			43.9*	30.2*	23*	38.7*	36.6*	35.9*
TTO	37.7	40.2	38.9	39.7	43	42.4	43.2	44.2
TUN	33.2	32.6	31.5	33	31.3	33.3	31.2	32.2
TUR	35.3	36.2	38.3	38.6	37.4	38.8	37.7	36.3
TZA	46.6	46.6	43.9	46.8	48.4	48.8	49.6	49.5
UGA	44.3	47.2	47.1	49	50.4	49.8	49.7	51.2
UKR	38.5	39.1	38.1	37.6	37.6	37.7	35.5	34.5
URY	47.8	50.2	48.7	48.7	50.1	49.9	49.1	49.8
USA	45.4	44	44.1	44.4	45.3	45.1	45.1	44.6
UZB	42.4	42	41.8	39.1	36.9	40.9	40.3	36.6
VEN [†]	51.2	49.3	49.8	50.9	51.5	50.2		
VNM	53	49.8	44.4	47	49.2	49.4	51	52.1
YEM [†]	51	55.4	55	55.3	54.1	57.3	59.2	
ZAF	66.7	67.4	69.2	69.4	68.7	67.9	65.6	62.2
ZMB	35.6*	44	45.1	41.3*	49.2	50.2	52.3	55.4
ZWE	35.1	39.9*	29.6*	42.8	38.7*	43.6	43.7	44.9

Notes: Table presents estimates of net income Gini index. Estimates missing if: [◊] trade data missing, [†] GNI per capita data missing, * estimated coefficient of variation squared negative. Values marked with * if deemed too sensitive to specific products imported.

F Data sources

Table XV: Original inequality data sets contained in WIID

Original source in WIID - Selection for first stage	Quality rating	Number of observations
Luxembourg Income Study (LIS) 2019	1	213
ECLAC 2019	1	59
Eurostat 2019	1	47
OECD 2018	1	47
OECD 2019	1	33
ECLAC	2	12
UNICEF 2004	2	7
Cambodia National Institute of Statistics	1	3
Commitment to Equity Project	1	3
UNICEF 2004	1	3
Cheong 2005	1	2
Chuliang et al. 2018	1	2
Deiningen and Squire, World Bank 2004	2	2
Ministry of Social Development	1	2
Szekely 2003	1	2
UNICEF 2007	2	2
Chotikapanich et al. 2005	2	1
Deiningen and Squire, World Bank 2004	1	1
Hong Kong Census and Statistics Department	1	1
Leibbrandt et al. 2009	2	1
Leibbrandt et al. 2010	1	1
Leibbrandt et al. 2010	2	1
Milanovic 1998	1	1
National Bureau of Statistics of China	1	1
UNICEF 2005	2	1
UNICEF 2011	1	1
UNICEF 2011	2	1
Whiteford and Van Seventer 2000	2	1

Notes: Table lists original data source and quality as given in WIID for Gini observations used in first stage. Quality ratings: 1 = High, 2 = Average. For details see User Guide WIID.

Table XVI: Data sources

Source	Variable	Notes
CEPII/BACI	trade flows	HS 6 digit, aggregated to 4 digit. Consumption goods classified by UNCTAD. Further excluded are HS chapters 26-27 “Mineral Products”, chapters 31, 34, 35, 36, 38 “Products of Chemical or Allied Industries” as well as chapters 72-81 and chapters 83, 84 of “Base Metals and Articles of Base Metal” and finally products listed under subheading 93 “Arms and Ammunition”.
CEPII/Gravity	distance, colonial dummy, common official language dummy, shared border dummy	
M. Larch’s Regional Trade Agreements Database from Egger and Larch (2008)	RTA dummy	
World Development Indicators	GNI per capita (PPP)	
UNU-WIDER: WIID	Gini	used in first stage. Welfare definition: “Income, disposable”; quality: “Average” & “High”; age coverage: “All”; area coverage: “All”; income unit: “Household”; unit of analysis: “Person”. Further selection among duplicate entries is based on sources: LIS 2019, OECD, SEDLAC/ECLAC, Eurostat. Additionally, non-duplicate entries fulfilling the unit selection criteria listed above are included.
LIS	Gini	for comparison, accessed March 2021
Povcal	Gini	for comparison, accessed March 2021
World Inequality Database	Gini	for comparison, accessed March 2021