

INTERNATIONAL TRADE PUZZLES: A SOLUTION LINKING PRODUCTION AND PREFERENCES

Justin Caron, Thibault Fally and James R. Markusen*

March 2014

Forthcoming in the Quarterly Journal of Economics

Abstract

International trade literature tends to focus heavily on the production side of general equilibrium, leaving us with a number of empirical puzzles. There is, for example, considerably less world trade than predicted by Heckscher-Ohlin-Vanek (HOV) models. Trade among rich countries is higher and trade between rich and poor countries lower than suggested by HOV and other supply-driven theories, and trade-to-GDP ratios are higher in rich countries. Our approach focuses on the relationship between characteristics of goods and services in production and characteristics of preferences. In particular, we find a strong and significant positive correlation of more than 45% between a good's skilled-labor intensity and its income elasticity, even when accounting for trade costs and cross-country price differences. Exploring the implications of this correlation for empirical trade puzzles, we find that it can reduce HOV's over-prediction of the variance of the net factor content of trade relative to that in the data by about 60%. Since rich countries are relatively skilled-labor abundant, they are relatively specialized in consuming the same goods and services that they are specialized in producing, and so trade more with one another than with poor countries. We also find a positive sector-level correlation between income elasticity and a sector's tradability, which helps explain the higher trade-to-GDP ratios in high-income relative to low-income countries.

Keywords: Non-homothetic preferences, gravity, income, missing trade.

JEL Classification: F10, F16, O10.

*Justin Caron: MIT Joint Program on the Science and Policy of Global Change and ETH Zurich; Thibault Fally: University of Colorado-Boulder (at the beginning of the project) and University of California-Berkeley; James Markusen: University of Colorado-Boulder. We thank Donald Davis, Peter Egger, Ana-Cecilia Fieler, Lionel Fontanie, Juan Carlos Hallak, Gordon Hanson, Jerry Hausman, Elhanan Helpman, Larry Karp, Wolfgang Keller, Ethan Ligon, Keith Maskus, Tobias Seidel, Ina Simonovska, David Weinstein, anonymous referees, as well as conference and seminar participants at the NBER Summer Institute (ITI), Society of Economic Dynamics, ERWIT CEPR conference, AEA-ASSA Meetings, Midwest Trade Meetings, UC San Diego, Paris School of Economics, Singapore Management University, ETH Zurich and the University of Colorado-Boulder for helpful comments. Contact: Thibault Fally, Department of Agricultural and Resource Economics, Giannini Hall, University of California, Berkeley CA 94720-3310; phone: 720-237-6158; fally@berkeley.edu.

1 Introduction

International trade theory is a general-equilibrium discipline. Yet it is probably fair to suggest that most of the standard portfolio of research focuses on the production side of general equilibrium. Price elasticities of demand do play a role in oligopoly models and, of course, a preference for diversity is important in all models, not just monopolistic competition. Income elasticities of demand are, however, generally assumed to be either one (homothetic preferences) or zero (so-called quasi-linear preferences used in oligopoly models). While non-homothetic preferences and the role of non-unitary income elasticities were crucial in the work of Linder (1961), subsequent work was limited. More recently, we see renewed interest in several strands of literature, including an important one on product quality.

These recent advances notwithstanding, we have a limited set of theoretical and empirical results regarding possible relationships between the demand and supply sides of general equilibrium; that is, not much is understood about whether certain characteristics of goods in production are correlated with other characteristics of preferences and demand. The purpose of our paper is to investigate such a relationship empirically. In particular, we explore a systematic relationship between factor intensities of goods in production and their corresponding income elasticities of demand in consumption. The existence of such a relationship can contribute to a number of empirical puzzles in trade as suggested by Markusen (2013). These include: i) the mystery of the missing net factor content of trade, ii) a home bias in consumption and iii) large trade volumes among rich countries and small trade volumes between rich and poor countries.

Our first objective is to estimate the importance of per capita income in determining demand patterns. Our results are derived from what we will term “constant relative income elasticity” (CRIE) preferences, recently used in Fieler (2011).¹ These are integrated within a general equilibrium model whose supply-side structure is based on an extension of Costinot et al. (2012) and Eaton and Kortum (2002) with multiple factors of production and an input-output structure as in Caliendo and Parro (2012). One immediate difficulty we face in the estimation of preferences is that we have expenditure data, not separate price and quantity data. This is a problem since trade costs can imply that goods are relatively cheaper in the country where they are produced, so large expenditure shares on home-produced (comparative advantage) goods may be partly due to trade costs. We solve this problem with a two-step estimation strategy. First, we use gravity equations to estimate patterns of comparative advantage and trade costs, and show that these can be used to compute price index proxies. In the second step, we use these

¹We also provide a discussion of alternative representations of non-homothetic preferences and expressions for expenditure shares across goods: the linear expenditure system, derived from Stone-Geary preferences, and Deaton and Muellbauer’s almost ideal demand system (AIDS) (Deaton and Muellbauer, 1980).

indexes to structurally control for price differences and estimate the parameters determining the income elasticity of demand. While the estimation of models with non-homothetic preferences has been considered as challenging in the past, our method is actually quite simple to implement as it does not rely on actual price data.² It is inspired from Redding and Venables (2004) and would also be consistent with a monopolistic-competition framework yielding gravity equations within each sector, as in Redding and Venables (2004) and Chaney (2008).³

Our estimations rely on the Global Trade Analysis Project (GTAP) data set, which comprises 94 countries with a wide range of income levels, 56 broad sectors including manufacturing and services, and 5 factors of production including the disaggregation of skilled and unskilled labor. This is an excellent data set for our purposes since it includes harmonized production, input-output, expenditure and trade data. However, the broad categories of goods and services make it unsuitable for the discussion of issues related to product quality and within-industry heterogeneity.

Results show that the income elasticity of demand varies considerably across goods from different industries. Moreover, it is significantly related both in economic and statistical terms to the skill intensity of a sector, with a correlation of over 50%. This fact has not yet been documented in the literature.⁴ As expected, accounting for trade costs and supply-side characteristics reduces this correlation, but it remains large and highly statistically significant. The relationship to capital intensity is positive but much weaker in economic terms and not statistically significant, consistent with Reimer and Hertel (2010), while the correlation with natural-resource intensity is negative.

The estimated parameters are then used to assess the role of per-capita income and non-homothetic preferences in explaining the empirical trade puzzles mentioned above. In addition to the income-elasticity / factor-intensity relationship, results include the following. First, a systematic relationship between income elasticity and skill intensity at the sector level generates a strong correlation between the factor content of production and consumption across countries. While about half of this correlation can be explained by trade costs, we find non-homotheticity to be as important quantitatively. This systematic relationship also contributes to solving a large part of the “missing trade” puzzle (Trefler, 1995). Standard Heckscher-Ohlin-Vanek Models with homothetic preferences famously predict a much larger variance in the net factor content of trade than what is seen in the data. We find that non-homothetic preferences reduce this excess variance by 60%, even after accounting for trade costs and factors embodied in

²As a robustness check, we use actual price data from the International Comparison Program (ICP).

³While the two-step estimation is more robust to misspecifications, we also propose a one-step estimation imposing additional restrictions on the demand and supply sides.

⁴A similar relationship has been emphasized by Verhoogen (2008) regarding quality: the production of high-quality goods tends to involve skilled workers.

intermediate goods.

Second, we illustrate how per-capita income and non-homothetic preferences can help us better understand patterns of bilateral trade volumes, in particular the low share of trade between rich and poor countries (North-South trade). Since high-income countries tend to be relatively abundant in skilled labor, the correlation between income elasticity and skill intensity implies that richer countries tend to consume goods for which they have a comparative advantage. Hence, they tend to trade more with one another than with low-income countries. As this mechanism also explains why countries will source a larger share of their consumption from themselves, non-homotheticity also contributes to explaining why, apart from trade costs, aggregate trade-to-GDP ratios are not higher than they are (the “home bias puzzle”). Furthermore, we identify a positive sector-level correlation between income elasticity and a sector’s tradability, which contributes to explaining why rich countries have higher trade-to-GDP ratios than low-income countries. Overall, allowing for non-homotheticity largely improves our understanding of the relationship between per-capita income and openness to trade, and explains part of the “home bias puzzle” for developing countries.

Our paper mainly contributes to two branches of the literature, one focusing on trade volumes and the other focusing on the factor content of trade. Early papers exploring the relationship between trade volumes and income elasticities are Markusen (1986), Hunter and Markusen (1988), Hunter (1991) and Bergstrand (1990). A particular focus of this literature is on the volume of trade in aggregate and among sets of countries, and its relationship to a world of identical and homothetic preferences as generally assumed in traditional trade theory. A general conclusion of this research is that non-homotheticity reduces trade volumes among countries with different per-capita income levels, though trade among high-income countries can increase. Matsuyama (2000) uses a competitive Ricardian model to arrive at a similar prediction. There has been a renewed interest in the role of preferences in explaining trade volumes recently, including Fieler (2011), Bernasconi (2011), Martinez-Zarzoso and Vollmer (2011) and Simonovska (2010).

Closest to our paper is Fieler (2011). She shows that aggregate trade data between countries are consistent with a model with two types of goods in which rich countries have a comparative advantage in income-elastic goods. This mechanism generates smaller trade flows between rich and poor countries. In Fieler (2011), income-elastic goods are also characterized by a higher dispersion of productivity and a lower elasticity of trade to trade costs, which can explain lower trade-to-GDP ratios for poorer countries. In our paper, we instead examine sector-level data and find evidence of a strong correlation between income elasticity and skill intensity rather than productivity dispersion,⁵ and that this correlation can better explain the lack of trade

⁵Note that we estimate the productivity dispersion parameter θ_k by sector while Fieler (2011) estimates

between poor and rich countries. Another difference is that Fielers (2011) considers only one factor of production (labor).

Another distinct branch of the literature to which we contribute has examined the net factor content of trade and the predictions of the Heckscher-Ohlin-Vanek Model. As in Trefler (1995) and Davis and Weinstein (2001), most of the attention has been put on the home bias or trade costs.⁶ Recent papers, including Cassing and Nishioka (2009) and Reimer and Hertel (2010), have emphasized the role of consumption patterns in explaining part of the “missing trade” puzzle but our results present several contributions. Cassing and Nishioka (2009) show that allowing for richer consumption patterns yields larger improvements in explaining the data than allowing for heterogeneous production techniques. They do not however specifically estimate non-homothetic preferences and cannot examine how much of the missing trade can actually be attributed to non-homotheticity. Both Cassing and Nishioka (2009) and Reimer and Hertel (2010) put an emphasis on capital intensity, which is positively but not strongly correlated with income elasticity of final demand, but they do not differentiate skilled from unskilled labor and thus underestimate the role of non-homothetic preferences in explaining the missing trade puzzle.

There are other topic areas where per-capita income plays a key role. One is a large and growing literature on product quality where per-capita income clearly matters: if a consumer is to buy one unit of a good, consumers with higher incomes buy higher quality goods. In line with Linder (1961), the role of quality differentiation has been underscored by Hallak (2006), Hallak (2010), Khandelwal (2010), Hallak and Schott (2011) and Fajgelbaum et al. (2011) among others. In addition, the distribution of income within a country matters, and a fairly general result is that higher inequality leads to a higher aggregate demand for high-quality products. We view this literature as important and most welcome. Note that within-industry reallocations only reinforce the mechanisms described in our model. If high-quality goods are associated with both higher income elasticities and stronger skill intensity, the same mechanisms would apply for within-industry reallocations as for the between-industry reallocations described in our paper.⁷

The rest of the paper is organized in three sections. We describe our theoretical framework

it for two broad categories of goods which could be the aggregation of various types of goods in our model. Simonovska and Waugh (2010) also estimate an aggregate productivity dispersion parameter but find little evidence that it differs between rich and poor countries.

⁶Here, we directly estimate the border effect, or equivalently a home bias in consumption, in the first-step gravity equation for each industry and control for it when we compare homothetic and non-homothetic preferences.

⁷Accounting for within-country inequalities only strengthens our results. We find similar estimates and slightly more variability in income elasticities when using within-country income distribution data by decile (see Appendix F).

in Section 2, our empirical strategy and estimation results in Section 3, and the implications for trade patterns and trade puzzles in Section 4.

2 Theoretical framework

2.1 Benchmark Model set-up

Demand

The economy is constituted of heterogeneous industries. In turn, each industry k is composed of a continuum of product varieties indexed by $j_k \in [0, 1]$. Preferences take the form:

$$U = \sum_k \alpha_{1,k} Q_k^{\frac{\sigma_k - 1}{\sigma_k}}$$

where $\alpha_{1,k}$ is a constant (for each industry k) and Q_k is a CES aggregate:

$$Q_k = \left(\int_{j_k=0}^1 q(j_k)^{\frac{\xi_k - 1}{\xi_k}} dj_k \right)^{\frac{\xi_k}{\xi_k - 1}}$$

Preferences are identical across countries, but non-homothetic if σ_k varies across industries. If $\sigma_k = \sigma$, we are back to traditional homothetic CES preferences. These preferences are used in Fieler (2011), with early analyses and applications found in Hanoch (1975) and Chao and Manne (1982). To the best of our knowledge, there is no common name attached to these preferences, so we will refer to them as constant relative income elasticity (CRIE) tastes: As shown in Fieler (2011) and below, the ratio of income elasticities of demand between goods i and j is given by σ_i/σ_j and is constant.

The CES price index of goods from industry k in country n is $P_{nk} = \left(\int_0^1 p_{nk}(j_k)^{1-\xi_k} dj_k \right)^{\frac{1}{1-\xi_k}}$. Given this price index, individual expenditures ($P_{nk}Q_{nk}$) in country n for goods in industry k equal:

$$x_{nk} = \lambda_n^{-\sigma_k} \alpha_{2,k} (P_{nk})^{1-\sigma_k} \quad (1)$$

where λ_n is the Lagrangian multiplier associated with the budget constraint of individuals in country n , and $\alpha_{2,k} = (\alpha_{1,k} \frac{\sigma_k - 1}{\sigma_k})^{\sigma_k}$. The Lagrangian λ_n is determined by the budget constraint: total expenditures must equal total income. In general there is no analytical expression for λ_n .

The income elasticity of demand η_{nk} for goods in industry k and country n equals:

$$\eta_{nk} = \sigma_k \cdot \frac{\sum_{k'} x_{nk'}}{\sum_{k'} \sigma_{k'} x_{nk'}} \quad (2)$$

It is clear from Equation (2) that the ratio of the income elasticities of any pair of goods k and k' equals the ratio of their σ parameters: $\frac{\eta_{nk}}{\eta_{nk'}} = \frac{\sigma_k}{\sigma_{k'}}$ and is constant across countries. Note that CRIE preferences (and separable preferences in general) preclude any inferior good: the income elasticity of demand is always positive for any good.⁸

Production

We assume Cobb-Douglas production functions with constant returns to scale: production depends on factors and bundles of intermediate goods from each industry. We assume that factors of production are perfectly mobile across sectors but immobile across countries. We denote by w_{fn} the price of factor f in country n . Factor intensities for each industry k and factor f are denoted by β_{kf} . We denote by γ_{kh} the share of the input bundles from industry h in total costs of industry k (direct input-output coefficient), and each input bundle is a CES aggregate of all varieties available in this industry (for the sake of exposition we assume that the elasticity of substitution between varieties is the same as for final goods). Total factor productivity $Z_{ik}(j_k)$ varies by country, industry and variety.

As common in the trade literature, we assume iceberg transport costs $d_{nik} \geq 1$ from country i to country n in sector k . The unit cost of supplying variety j_k to country n from country i equals:

$$p_{nik}(j_k) = \frac{d_{nik}}{Z_{ik}(j_k)} \prod_f (w_{fi})^{\beta_{kf}} \prod_h (P_{hi})^{\gamma_{kh}}$$

where P_{ih} is the price index of goods h in country i and $\sum_f \beta_{kf} + \sum_h \gamma_{kh} = 1$.

There is perfect competition for the supply of each variety j_k . Hence, the price of variety j_k in country n in industry k equals:

$$p_{nk}(j_k) = \min_i \{p_{nik}(j_k)\}$$

We follow Eaton and Kortum (2002) and assume that productivity $Z_{ik}(j_k)$ is a random variable with a Frechet distribution. This setting generates gravity within each sector. Productivity is independently drawn in each country i and industry k , with a cumulative distribution:

$$F_{ik}(z) = \exp \left[-(z/z_{ik})^{-\theta_k} \right]$$

where z_{ik} is a productivity shifter reflecting average TFP of country i in sector k . As in Eaton and Kortum (2002), θ_k is related to the inverse of productivity dispersion across varieties within

⁸Another notable feature of income elasticities is that they decrease with income. A larger income induces a larger fraction of expenditures in high- σ_k industries. Hence, the consumption-weighted average of σ_k is larger (denominator in expression 2 above), which yields lower income elasticities.

each sector k . Note that we also assume $\theta_k > \xi_k - 1$ to ensure a well-defined CES price index within each industry.

In the benchmark version of the model, we allow the dispersion parameter θ_k to vary across industries. As in Costinot, Donaldson and Komunjer (2010), we also allow the shift parameter z_{ik} to vary across exporters and industries, keeping a flexible structure on the supply side and controlling for any pattern of Ricardian comparative advantage forces at the sector level.

Endowments

Each country i is populated by a number L_i of individuals. The total supply of factor f is fixed in each country and denoted by V_{if} . As a first approximation, each person is endowed by V_{if}/L_i units of factor V_{fi} implying no within-country income inequality. We relax this assumption in Appendix F and examine how within-country income inequality affects our estimates.

2.2 Two special cases

Our benchmark specification of the supply-side is very flexible and allows for several sources of comparative advantage. We also propose two more restrictive alternative production specifications to better illustrate the interaction between supply-side characteristics and non-homotheticity on the demand side.

- **Skill-driven Model:**

In this special case, we impose the dispersion of productivity $1/\theta$ to be equal across all sectors. We also impose the same productivity shifters across all sectors in each country:

Additional assumptions in the Skill-driven Model:

- i) Fréchet dispersion parameters $\theta_k = \theta$ are constant across sectors.*
- ii) Productivity shifters $z_{ik} = z_i$ are constant across all sectors for each exporter i .*

In this more restrictive version of the model, forces of Ricardian comparative advantage are assumed away. With common θ 's and common z 's across sectors, the distribution of TFP is the same across sectors (for a given country). To better illustrate the role of skill intensity, we also assume that there are only two factors of production: unskilled labor and skilled labor, and no intermediate goods.

- **Theta-driven Model (Fieler, 2011):**

This particular case replicates across industries the assumptions made in Fieler (2011) across (unobserved) types of goods. It allows for variations in θ_k across industries but assumes that the technology parameter T_i , defined here as $T_i \equiv z_{ik}^{\theta_k}$, is constant across industries. Moreover, as in Fieler (2011), it only considers one factor of production and neglects the differences in factor endowments:

Additional assumptions in the Theta-driven Model:

- i) Productivity shifters are given by $z_{ik} = T_i^{1/\theta_k}$ where T_i is constant across sectors for each exporter i .*
- ii) Labor is only one factor of production $f = L$.*

These assumptions generate a comparative advantage for rich countries (high- T countries) in low- θ sectors, i.e. sectors with more dispersed productivity. Following Costinot, Donaldson and Komunjer (2012), this result derives from the ranking in relative productivity shifters:

$$\frac{z_{ik}}{z_{i'k}} = \left(\frac{T_i}{T_{i'}}\right)^{1/\theta_k} > \left(\frac{T_i}{T_{i'}}\right)^{1/\theta_{k'}} = \frac{z_{ik'}}{z_{i'k'}}$$

if $T_i > T_{i'}$ and $\theta_k < \theta_{k'}$. Since θ_k also governs the elasticity of trade to trade costs, Fieler (2011) imposes rich countries to have a comparative advantage in goods for which there are higher incentives to trade.

For all versions of the model, equilibrium is characterized by the same set of market conditions described in the next subsection.

2.3 Equilibrium

Equilibrium is defined by the following equations. On the demand side, total expenditures D_{nk} of country n in final goods k simply equals population L_n times individual expenditures as shown in (1). This gives:

$$D_{nk} = L_n (\lambda_n)^{-\sigma_k} \alpha_{2,k} (P_{nk})^{1-\sigma_k} \quad (3)$$

where $\alpha_{2,k}$ is an industry constant defined in equation (1). λ_n is the Lagrangian multiplier associated with the budget constraint:

$$L_n e_n = \sum_k D_{nk} \quad (4)$$

where e_n denotes per-capita income. Total demand X_{nk} for goods k in country n is the sum of the demand for final consumption D_{nk} and intermediate use:

$$X_{nk} = D_{nk} + \sum_h \gamma_{kh} Y_{nh} \quad (5)$$

where Y_{nh} refers to total production in sector h .

On the supply side, each industry mimics an Eaton and Kortum (2002) economy. In particular, given the Frechet distribution, we obtain a gravity equation for each industry. We follow Eaton and Kortum (2002) notation with the addition of industry subscripts. By denoting X_{nik} the value of trade *from* country i to country n , we obtain:

$$X_{nik} = \frac{S_{ik}(d_{nik})^{-\theta_k}}{\Phi_{nk}} X_{nk} \quad (6)$$

where S_{ik} and Φ_{nk} are defined as follows. The “supplier effect”, S_{ik} , is inversely related to the cost of production in country i and industry k . It depends on the total factor productivity parameter z_{ik} , intermediate goods and factor prices:

$$S_{ik} = z_{ik}^{\theta_k} \left(\prod_f (w_{fi})^{\beta_{kf}} \right)^{-\theta_k} \left(\prod_h (P_{ih})^{\gamma_{kh}} \right)^{-\theta_k} \quad (7)$$

The parameter θ_k is inversely related to the dispersion of productivity within sectors, implying that differences in productivity and factor prices across countries have a stronger impact on trade flows in sectors with higher θ_k .

In turn, we define Φ_{nk} as the sum of exporter fixed effects deflated by trade costs. Φ_{nk} plays the same role as the “inward multilateral trade resistance index” as in Anderson and van Wincoop (2003):

$$\Phi_{nk} = \sum_i S_{ik}(d_{nik})^{-\theta_k} \quad (8)$$

This Φ_{nk} is actually closely related to the price index, as in Eaton and Kortum (2002):

$$P_{nk} = \alpha_{3,k} (\Phi_{nk})^{-\frac{1}{\theta_k}} \quad (9)$$

with $\alpha_{3,k} = \left[\Gamma \left(\frac{\theta_k + 1 - \xi_k}{\theta_k} \right) \right]^{\frac{1}{\xi_k - 1}}$ where Γ denotes the gamma function.⁹

⁹Alternatively, we can generalize this model and assume that the elasticity of substitution for intermediate use differs from the elasticity of substitution for final use, and depends on the parent industry. This does not affect the elasticity of the price index w.r.t. Φ_k as long as the dispersion parameter θ_k does not depend on the final use. Differences in elasticities of substitution would then be captured by the industry fixed effect that we include in our estimation strategy and would not affect our estimates.

Finally, two other market clearing conditions are required to determine factor prices and income in general equilibrium. Given the Cobb-Douglas production function, total income from a particular factor equals the sum of total production weighted by the factor intensity coefficient β_{kf} . With factor supply V_{fi} and factor price w_{fi} for factor f in country i , factor market clearing implies:

$$V_{fi}w_{fi} = \sum_k \beta_{kf}Y_{ik} \quad (10)$$

where output equals the sum of outward flows $Y_{ik} = \sum_n X_{nik}$. In turn, per-capita income is determined by:

$$e_i = \frac{1}{L_i} \sum_f V_{fi}w_{fi} \quad (11)$$

By Walras' Law, trade is balanced at equilibrium.

The two special cases of the benchmark model share the same set of equilibrium conditions. Note that S_{ik} takes a more specific form in each case:

- Skill-driven Model:

$$S_{ik} = z_i^\theta \left(\prod_f (w_{fi})^{\beta_{kf}} \right)^{-\theta} \quad (12)$$

- Theta-driven Model:

$$S_{ik} = T_i w_i^{-\theta_k} \quad (13)$$

2.4 Implications: the role of non-homothetic preferences

2.4.1 Trade partners and volumes of North-South Trade

With non-homothetic preferences, differences in income per capita across countries can result in large differences in consumption patterns, even though preferences are assumed identical. In this section, we illustrate how non-homotheticity affects trade patterns when there is a systematic relationship between preference parameters and characteristics of the supply side, e.g. factor intensities. Such a relationship is supported by our empirical analysis which finds, in particular, a positive correlation across sectors between skilled-labor intensity and income elasticity.

Let us first consider the case in which trade costs are assumed away ($d_{nik} = 1$). In this case, prices are the same in all countries and the share of consumption corresponding to imports from i in industry k is the same for all importers (country n): $\frac{X_{nik}}{D_{nk}} = \frac{S_{ik}}{\sum_j S_{jk}}$. Assuming no trade in intermediates and summing over all industries, total import penetration by country i

in country n is:

$$\frac{X_{ni}}{X_n} = \sum_k \left(\frac{S_{ik}}{\sum_j S_{jk}} \right) \left(\frac{\alpha_{4,k} \lambda_n^{-\sigma_k}}{\sum_{k'} \alpha_{4,k'} \lambda_n^{-\sigma_{k'}}} \right) \quad (14)$$

where $X_n = L_n e_n$ is total expenditures in country n , $X_{ni} = \sum_k X_{nik}$ is total bilateral trade from country i to n , and $\alpha_{4,k} = \alpha_{2,k} P_k^{1-\sigma_k}$ is an industry constant incorporating common prices. The first term in parentheses is the share of imports from i in consumption of k – it reflects the comparative advantage of country i in sector k . The second term corresponds to the share of industry k in final consumption of country n .

Aggregate import penetration by country i in country n obviously depends on the sectoral composition of both supply and demand, but the latter has generally been neglected by previous work. If preferences are homothetic, $\sigma_k = \sigma$ is common across industries and import penetration is the same across all importers n (for a given exporter i). When preferences are non-homothetic (heterogenous σ_k), exporters with a comparative advantage in high- σ industries have a relatively larger penetration in rich countries (low λ_n), while exporters with a comparative advantage in low- σ industries have a relatively larger penetration in poor countries (high λ_n). We will show empirically that rich countries have a comparative advantage in high- σ (also skill-intensive) industries which can quantitatively explain large differences in trade volumes across country pairs depending on each partner's per-capita income.¹⁰

Trade costs provide an alternative explanation as to why import penetration varies across markets. On the supply side, proximity reduces unit costs. On the demand side, consumption might be biased towards goods produced locally if their price is lower (e.g. Saudi Arabia consuming relatively more petroleum). The latter argument requires that the elasticity of substitution be larger than one. These effects of trade costs can reinforce the patterns described above. In our framework, a general expression for the import penetration of exporter i in market n yields:

$$\frac{X_{ni}}{X_n} = \sum_k \pi_{nik} s h_{nk} = \sum_k \left(\frac{S_{ik} d_{nik}^{-\theta_k}}{\Phi_{nk}} \right) \left(\frac{\alpha_{5,k} \lambda_n^{-\sigma_k} \Phi_{nk}^{\frac{\sigma_k-1}{\theta_k}}}{\sum_{k'} \alpha_{5,k'} \lambda_n^{-\sigma_{k'}} \Phi_{nk'}^{\frac{\sigma_{k'}-1}{\theta_{k'}}}} \right) \quad (15)$$

where $\Phi_{nk} = \sum_j S_{jk} d_{nj}^{-\theta_k}$ by definition (Equation 8) and $\alpha_{5,k} = \alpha_{2,k} \alpha_{3,k}^{1-\sigma_k}$ is an industry constant. The first term in parentheses corresponds to π_{nik} , the share of imports in n from country i in sector k while the second term corresponds to $s h_{nk}$, the share of sector k in consumption in country n . Import shares and consumption shares are both affected by trade costs. In the

¹⁰Formally, if per capita income e_n increases with n , if S_{ik} is log-supermodular (i.e. countries with higher index i have a comparative advantage in sectors with higher index k as in Costinot (2009)), and if σ_k increases with k , then X_{ni} is log-supermodular, which means that $\frac{X_{ni}}{X_{ni'}} > \frac{X_{n'i}}{X_{n'i'}}$ for any countries $n > n'$ and $i > i'$. The proof follows from Athey (2002) since both S_{ik} and $\lambda_n^{-\sigma_k}$ are log-supermodular.

empirical section, we thus need to carefully examine the distinct contribution of trade costs and non-homotheticity. In addition, we should note that import penetration by exporter i in rich countries might not increase with exporter i 's per capita income if competition effects dominate demand effects.¹¹ For instance, a car producer may find it difficult to export cars to Germany because of trade costs and competition with local producers, even if Germany has a relatively large consumption of cars. Our empirical results however indicate that demand effects dominate.

2.4.2 Openness and the home bias puzzle

Non-homothetic preferences can also influence aggregate trade-to-GDP ratios – a key measure often used as an indicator of a country's openness to trade – through at least two channels. First, if high-income countries tend to have a comparative advantage in income-elastic goods, countries at either end of the income distribution will consume larger shares of their own goods than would be predicted under homothetic preferences. This induces a lower trade-to-GDP ratio and contributes to explaining the “home bias puzzle”.¹² Second, if trade costs are larger for low income-elasticity goods, or if trade is more sensitive to trade costs for such goods (as in Fieler, 2011), observed aggregate openness will tend to be lower for poorer countries.

To illustrate these two channels, let us examine $\bar{\pi}_{nn}$ the aggregate share of goods purchased internally in country n (equal to one minus the aggregate share of imports over total demand). It equals:

$$\bar{\pi}_{nn} = \sum_k \pi_{nnk} sh_{nk}$$

where $\pi_{nnk} \equiv \frac{X_{nnk}}{X_{nk}} = \frac{S_{nk}d_{nnk}^{-\theta_k}}{\sum_i S_{ik}d_{nik}^{-\theta_k}}$ is the share of good k purchased from domestic production. This share only depends on supply-side characteristics (trade costs and the relative cost of producing goods k). The second term $sh_{nk} \equiv \frac{X_{nk}}{\sum_{k'} X_{nk'}}$ is the share of good k in total demand in country n . Let us also denote by $\pi_{av,k} = \frac{1}{N} \sum_n \pi_{nnk}$ the average share of goods purchased internally in sector k (inversely related to good k 's tradability).¹³ Holding imports shares π_{nik} constant within each industry, we can examine the difference in the aggregate demand for domestic goods implied by differences between consumption patterns predicted by homothetic

¹¹Formally, this can arise when $\lambda_n^{-\sigma_k} \Phi_{nk}^{\frac{\sigma_k-1}{\theta_k}-1}$ is not log-supermodular, even if $\lambda_n^{-\sigma_k}$ is log-supermodular.

¹²Low levels of international-to-domestic trade flows have been discussed by McCallum (1995), Anderson and van Wincoop (2003), Yi (2010), among others.

¹³Note we would have $\pi_{av,k} = \frac{1}{N}$ be the same across all goods and countries if there were no trade costs.

and non-homothetic preferences:

$$\bar{\pi}_{nn}^{NH} - \bar{\pi}_{nn}^H = \underbrace{\sum_k (\pi_{nnk} - \pi_{av,k}) (sh_{nk}^{NH} - sh_{nk}^H)}_{\text{covariance}} + \underbrace{\sum_k \pi_{av,k} (sh_{nk}^{NH} - sh_{nk}^H)}_{\text{tradability}} \quad (16)$$

where sh_{nk}^{NH} and sh_{nk}^H denote consumption shares for non-homothetic and homothetic preferences.

If non-homothetic preferences shift consumption towards goods in which countries have a comparative advantage (e.g. unskilled-labor intensive sectors in low-income countries), the first “covariance” term $\sum_k (\pi_{nnk} - \pi_{av,k}) (sh_{nk}^{NH} - sh_{nk}^H)$ should be positive on average, leading to lower aggregate demand for imported goods. This illustrates the first channel.

The second term of Equation (16) reflects the second mechanism described above. If income elasticities are systematically correlated with tradability $\pi_{av,k}$ across sectors, the second term $\sum_k \pi_{av,k} (sh_{nk}^{NH} - sh_{nk}^H)$ can be of a different sign in poor and rich countries. If trade costs are smaller for income-elastic goods, or if the elasticity of trade to trade costs is smaller for income-elastic goods (as in Fieler 2011), rich countries will tend to consume goods with smaller $\pi_{av,k}$. This generates larger trade-to-GDP ratios for rich countries than for poor countries.

We show in the empirical section that non-homothetic preferences play a role through both of these channels. Note that these mechanisms reinforce the effect of trade costs. In particular, an alternative explanation for low trade-to-GDP in poor countries is that trade costs are systematically higher in those countries (Waugh 2010). We illustrate quantitatively the role of non-homothetic preferences after accounting for trade costs in gravity equations for each sector.

2.4.3 Missing factor content of trade

One reason why comparative advantage may be related to consumption patterns is that the income elasticity of demand is correlated with skilled labor requirements. This provides rich countries, which are abundant in skilled labor, a comparative advantage in goods that rich consumers are more likely to buy. As we describe now, such a correlation can also shed light on the “missing trade puzzle” – the fact that the variance in the embodied factor content in net trade is so low relative to the predictions of the Heckscher-Ohlin-Vanek (HOV) model (Trefler, 1995).

Standard HOV models assume homothetic preferences. This assumption implies that, under costless trade, consumption shares for each industry are the same in all countries. Accounting for non-homothetic preferences can yield very different predictions in terms of factor content of trade. In particular, it can potentially explain why poor countries trade so little with rich

countries (in factor content) even if their endowments differ largely. The intuition is simple. When the income elasticity of demand is correlated with skill intensity, consumption in rich countries is biased towards skill-intensive industries, which also means that they are more likely to import from skill-abundant countries, i.e. rich countries. The same intuition would apply to capital if the income elasticity of demand would be correlated with capital intensity and if richer countries were relatively more endowed in capital.

This intuition can be simply illustrated in our framework. We define the factor content of trade F_{fn} as the *value* of factor f required to produce exports minus imports. It equals $F_{fn} = \sum_k \beta_{kf} (\sum_{i \neq n} X_{nik} - \sum_{i \neq n} X_{ink})$ when there is no intermediate goods trade and production coefficients β_{kf} are common across countries.¹⁴ After simple reformulations, we can decompose F_{fn} in two terms:

$$F_{fn} = s_n \sum_k \bar{Y}_k \beta_{kf} \left[\frac{Y_{nk}}{s_n \bar{Y}_k} - 1 \right] - s_n \sum_k \bar{Y}_k \beta_{kf} \left[\frac{D_{nk}}{s_n \bar{Y}_k} - 1 \right] \quad (17)$$

$$= \underbrace{s_n \sum_k \bar{Y}_k \beta_{kf} \left[\frac{Y_{nk}}{s_n \bar{Y}_k} - 1 \right]}_{F_{fn}^{HOV}} - \underbrace{s_n \sum_k \bar{Y}_k \beta_{kf} \left[\frac{D_{nk}}{s_n \bar{Y}_k} - 1 \right]}_{F_{fn}^{CB}} \quad (18)$$

where $Y_{nk} = \sum_i X_{ink}$ denotes the value of production of country n in sector k , $\bar{Y}_k = \sum_n Y_{nk}$ denotes the value of world's production in sector k , and s_n denotes the share of country n in world GDP. Note that we define factor content in terms of factor reward instead of quantities (number of workers or machines).¹⁵

In the brackets, the $\frac{D_{nk}}{s_n \bar{Y}_k}$ ratio equals the share of consumption of k in country n relative to the share of consumption of k in the world. The ratio $\frac{Y_{nk}}{s_n \bar{Y}_k}$ equals the share of production in sector k in country n relative to the share of production in sector k in the world. With homothetic preferences and costless trade, the second term in brackets would be null ($\frac{D_{nk}}{s_n \bar{Y}_k} - 1 = 0$) and the expression above could be simplified to:

$$F_{fn} = F_{fn}^{HOV} = w_{fn} V_{fn} - s_n \sum_i w_{fi} V_{fi} \quad (19)$$

Under factor price equalization w_{fn} is the same across countries and the above expression corresponds to the standard prediction of the net factor content trade in the HOV model. This

¹⁴The empirical section and the appendix derive additional results to account for traded intermediate inputs and production coefficients that differ across countries.

¹⁵Standard HOV estimation assumes factor price equalization. Under this assumption, both approaches are equivalent. When FPE is violated, for instance when factor productivity differs across countries, the predicted factor content has to be adjusted for such differences if written in terms of factor units (e.g. number of workers or machines). No adjustment is necessary if we focus on values, i.e. factor supply times factor prices. This approach simplifies the exposition of the main intuitions and better illustrates the contribution of non-homothetic preferences relative to homothetic preferences without providing too much detail on factor prices.

equation states that the amount of factor f embedded in country n 's exports should equal the total value of the supply of factor f in this country minus the value of the world's supply of this factor adjusted by the share s_n of country n in world GDP.

Equation (19) is violated when preferences are not homothetic and $\frac{D_{nk}}{s_n Y_k} - 1$ differs from zero. It thus needs to be corrected by a consumption term F_{fn}^{CB} (where ‘‘CB’’ stands for consumption bias). In particular, if relative consumption $\frac{D_{nk}}{s_n Y_k}$ is positively correlated with production $\frac{Y_{nk}}{s_n Y_k}$, then F_{fn}^{CB} is correlated with F_{fn}^{HOV} and predicted factor trade is smaller than predicted by models with homothetic preferences. In the empirical section, we verify that $\frac{D_{nk}}{s_n Y_k}$ and $\frac{Y_{nk}}{s_n Y_k}$ are indeed strongly correlated across countries and industries and that F_{fn}^{CB} is correlated with F_{fn}^{HOV} across countries and factors.

Again, trade costs contribute to the positive correlation between supply and demand across industries as well as the low factor content of trade as shown by Davis and Weinstein (2001). Additionally, asymmetric trade costs as in Waugh (2010) can explain low levels of trade to and from low-income countries and can potentially shed some light on the missing trade puzzle. In the empirical section, we disentangle the effect of trade costs and non-homothetic demand and show that the latter plays an important role. Also, differences in factor requirements across countries as well as trade in intermediate goods can also partially explain the missing trade puzzle. In the empirical section, we follow the methodology developed by Trefler and Zhu (2010) to illustrate the role of non-homotheticity, accounting for more complex vertical linkages.

3 Estimation

The first objective of this section is to detail the two-step estimation of the benchmark model as well as the one-step estimation of the two special cases (Skill-driven and Theta-driven Models), leading to the identification of the parameters determining the income elasticity of demand. The second objective is to test for a positive correlation between income elasticity and factor intensity.

3.1 Two-step estimation of the benchmark model

The value of final demand in an industry is determined as in Equation (3) or equivalently Equation (1) for individual expenditures $x_{nk} = \frac{D_{nk}}{L_n}$. In log, the model yields:

$$\log x_{nk} = -\sigma_k \cdot \log \lambda_n + \log \alpha_{2,k} + (1 - \sigma_k) \cdot \log P_{nk} \quad (20)$$

where $\alpha_{2,k}$ is a preference parameter which varies across industries only. In addition, final

demand should satisfy the budget constraint which determines λ_n : a higher income per capita is associated with a smaller Lagrangian multiplier λ_n .

If there were no trade costs, the price index P_{nk} would be the same across countries and could not be distinguished from an industry fixed effect. If, in richer countries, consumption were larger in a particular sector relative to other sectors, the estimated $\hat{\sigma}_k$ would be larger for this sector. Since trade is not costless, estimated income elasticities would be biased if we did not control for the price index P_{nk} (to capture supply-side characteristics). As richer countries have a comparative advantage in skill-intensive industries, the price index is relatively lower in these industries. Conversely, poor countries have a comparative advantage in unskilled labor intensive industries and thus have a lower price index in these industries relative to other industries. When the elasticity of substitution between industries is larger than one, these differences in price indexes in turn affect the patterns of consumption. If we were not controlling for P_{nk} , we would overestimate the income elasticity in skill intensive sectors.

We proceed in two steps. The main goal of the first step is to obtain a proxy for the price index $\log P_{nk}$. According to the equilibrium condition (9), $\log P_{nk}$ depends linearly on $\log \Phi_{nk}$ which can be identified using gravity equations. Then, using the estimated price indices (or equivalently $\hat{\Phi}_{nk}$), we can estimate the final demand equation (20).

Both steps follow the structure of the benchmark model (general case) and are also consistent with the two nested models (Skill-driven and Theta-driven). This two-step procedure estimates the supply-side and demand-side parameters separately, and is thus more robust to model misspecifications on either side. In Section 3.2, we also develop an alternative one-step estimation strategy to estimate the two Skill-driven and Theta-driven Models and exploit additional restrictions which affect both the supply and demand sides.

Step 1: Gravity equation estimation and identification of Φ_{nk}

By taking the log of trade flows in Equation (6), the model yields:

$$\log X_{nik} = \log S_{ik} - \theta_k \log d_{nik} + \log (X_{nk}/\Phi_{nk}) \quad (21)$$

We estimate this equation for each sector by including importer fixed effects (in place of $\log \frac{X_{nk}}{\Phi_{nk}}$) and exporter fixed effects (in place of $\log S_{ik}$) as well as proxies for trade costs d_{nik} . As we do not have data on bilateral transport costs by industry, we assume d_{nik} to be a log-linear combination of various trade cost variables:

$$\log d_{nik} = \sum_{var} \delta_{var,k} TC_{var,ni} + \delta_{ATC,ik} B_{i \neq n}$$

where $TC_{var,ni}$ refers to the variables (indexed by var) included in the gravity equation to capture trade costs between n and i . Following the literature on gravity, we include the log of physical distance (including internal distance), a common language dummy, a colonial link dummy, a border effect dummy (equal to one if $i \neq n$), a contiguity dummy (equal to one if countries i and n share a common border), a free-trade-agreement dummy (equal to one if there is an agreement between countries i and n), a common currency dummy and a common-legal-origin dummy (equal to one if i and n have the same legal origin: British, French, German, Scandinavian or socialist). Parameters $\delta_{var,k}$ capture the elasticity of trade costs to each trade cost variable var . They are indexed by k : the effect of each trade cost variable may differ across industries. Notice that all these proxies imply symmetric trade costs. Following Waugh (2010), we also consider asymmetric trade costs (ATC) by including exporter-specific border effects $\delta_{ATC,ik}B_{i \neq n}$ (where $B_{i \neq n}$ is a dummy equal to one for international trade flows and δ_{ik} is an exporter-specific coefficient).

Incorporating the expression for trade costs into the equation for trade flows (21), we obtain our estimated equation:

$$X_{nik} = \exp \left[FX_{ik} + FM_{nk} - \sum_{var} \beta_{var,k} TC_{var,ni} - \beta_{ATC,ik} B_{i \neq n} + \varepsilon_{nik}^G \right] \quad (22)$$

where ε_{nik}^G is the error term, FM_{nk} refers to importer fixed effects and FX_{ik} to exporter fixed effects, $\beta_{ATC,ik} = \theta_k \delta_{ATC,ik}$ and $\beta_{var,k} = \theta_k \delta_{var,k}$ for each trade cost variable var . Note that θ_k cannot be directly identified from $\delta_{var,k}$ using the gravity equation.¹⁶ Since all coefficients to be estimated are sector specific, we can estimate this gravity equation separately for each sector (as a result, we do not impose trade balance). Following Silva Santos and Tenreyro (2006), we estimate gravity using the Poisson pseudo-maximum likelihood estimator (Poisson PML).

The model tells us that importer and exporter fixed effects FX_{ik} and FM_{nk} capture valuable information on S_{ik} and Φ_{nk} . We follow a strategy developed by Redding and Venables (2004) to estimate Φ_{nk} .¹⁷ Following Equation (8) defining Φ_{nk} , we use the estimates of $\log S_{ik}$ (from \widehat{FX}_{ik}) and $\theta_k \log d_{nik}$ (from $\sum_{var} \widehat{\beta}_{var,k} TC_{var,ni}$) to construct a structural proxy for Φ_{nk} :

$$\widehat{\Phi}_{nk} = \sum_i \exp \left(\widehat{FX}_{ik} - \sum_{var} \widehat{\beta}_{var,k} TC_{var,ni} - \beta_{ATC,ik} B_{i \neq n} \right) \quad (23)$$

¹⁶In subsection 3.2 we examine an alternative specification (Theta-driven Model) where $\delta_{var,k}$ are assumed to be constant across sectors and where cross-sectoral variations in trade elasticities can be used to identify differences in θ_k across sectors.

¹⁷See also Fally et al. (2010), Head and Mayer (2006). An alternative method uses importer fixed effects and observed total demand to estimate Φ_{nk} . The two methods are actually equivalent when gravity is estimated with Poisson PML, see Fally (2012).

This constructed $\widehat{\Phi}_{nk}$ varies across industries and countries in an intuitive way. It is the sum of all potential exporters' fixed effect (reflecting unit costs of production) deflated by distance and other trade cost variables. If country n is close to an exporter that has a comparative advantage in industry k , i.e. an exporter associated with a large exporter fixed effect FX_{ik} (large S_{ik}), our constructed $\widehat{\Phi}_{nk}$ will be relatively larger for this country reflecting a lower price index of goods from industry k in country n . Note that $\widehat{\Phi}_{nk}$ also accounts for domestic supply in each industry k (when $i = n$).¹⁸

Such a method would fit various structural frameworks. If our model were based on a Dixit-Stiglitz-Krugman framework instead of Eaton-Kortum, price indices by importer and industry could be obtained in the same way. This could account for an endogenous range of available varieties.

Step 2: Demand system estimation and identification of σ_k

The first step estimation gives us an estimate of Φ_{nk} . From Equation (9), we know that the price index P_{nk} is a log-linear function of Φ_{nk} which we can use as a proxy for P_{nk} on the right-hand side of Equation (20) describing final demand.¹⁹ Our estimated equation for per capita final demand is thus:

$$\log x_{nk} = -\sigma_k \cdot \log \lambda_n + \log \alpha_{5,k} + \frac{(\sigma_k - 1)}{\theta_k} \log \widehat{\Phi}_{nk} + \varepsilon_{nk}^D \quad (24)$$

where ε_{nk}^D denotes the error term. In each country n , we further impose the sum of fitted expenditures across sectors to equal observed total per capita expenditures e_n :

$$\sum_k \exp \left[-\sigma_k \cdot \log \lambda_n + \log \alpha_{5,k} + \frac{(\sigma_k - 1)}{\theta_k} \log \widehat{\Phi}_{nk} \right] = e_n \quad (25)$$

We jointly estimate (24) and (25) using constrained non-linear least squares (we minimize the sum of squared errors $(\varepsilon_{nk}^D)^2$ while imposing both Equations 24 and 25 to hold). Observed variables are: the price proxies $\widehat{\Phi}_{nk}$, individual expenditures x_{nk} per industry and total expenditures e_n .²⁰ Free parameters to be estimated are the σ_k , the dispersion parameters θ_k , the Lagrangian multipliers λ_n and the industry fixed effects $\alpha_{5,k}$.

¹⁸Also note that the error term ε_{nik}^G is not included in the construction of $\widehat{\Phi}_{nk}$. An unobserved shock affecting trade for a specific country pair would not affect $\widehat{\Phi}_{nk}$. This mitigates potential omitted variable and endogeneity biases jointly affecting trade relationships and demand patterns.

¹⁹As a robustness check, we estimate the demand equation using actual price data instead or in addition to using $\log \widehat{\Phi}_{nk}$ (Appendix C).

²⁰Note that our data are micro-consistent. For each country, we have $\sum_k x_{nk} = e_n$

Two normalizations are required. Given the inclusion of industry fixed effects, λ_n can only be identified up to a constant.²¹ We thus normalize $\lambda_{USA} = 1$ for the US. A similar issue arises for σ_k which can be estimated only up to a common multiplier.²² We thus normalize $\sigma_{TEX} = 1$ for textiles. Despite this, income elasticities can be derived based on Equation (2):

$$\hat{\eta}_{nk} = \hat{\sigma}_k \cdot \frac{\sum_{k'} \hat{x}_{nk'}}{\sum_{k'} \hat{\sigma}_{k'} \hat{x}_{nk'}} \quad (26)$$

Multiplying all σ_k by the same constant has no effect on estimated income elasticities.

This estimation procedure can be seen as a non-linear least squares estimation of Equation (24) in which λ_n is the implicit solution of Equation (25) and thus a function of fitted coefficients and observed per capita expenditures e_n .²³ While this estimation procedure is consistent with general equilibrium conditions, we show that similar estimates are found when estimating Equation (24) either without constraining the sum of fitted expenditures to equal observed per capita expenditures e_n (Equation 25) or in a reduced-form approximation in which $\log \lambda_n$ is replaced by a linear function of $\log e_n$ (see Appendix B).

The benchmark specification described above identifies σ_k and income elasticities solely based on the coefficient associated with the Lagrangian λ_n . The σ_k parameter also appears in the coefficient for Φ_{nk} in Equation (24) but the benchmark specification does not impose any constraint on the coefficient for Φ_{nk} since θ_k is a free parameter (we can then identify θ_k using σ_k and the coefficient for Φ_{nk}). In an alternative estimation, we jointly identify σ_k from the coefficients on λ_n and Φ_{nk} by constraining θ_k to equal 4 in all sectors.²⁴ This choice of θ is close to the Simonovska and Waugh (2010) estimates of 4.12 and 4.03. Donaldson (forthcoming), Eaton et al. (2011), Costinot et al. (2012) provide alternative estimates that range between 3.6 and 5.2. Alternative values for θ yield very similar results for income elasticities.

Because $\hat{\Phi}_{nk}$ is a generated regressor, standard errors on the demand parameters must explicitly account for errors coming from the first-step estimations.²⁵ Because of the non-linearities arising in the computation of $\hat{\Phi}_{nk}$, we estimate bootstrap standard errors by resampling countries (importers) and sectors. For each bootstrap sample, we re-estimate the two steps: gravity

²¹To see this, we can multiply λ_n by a common multiplier λ' and multiply the industry fixed effect α_k by $(\lambda')^{\sigma_k}$. Using $\lambda_n \lambda'$ instead of λ_n and $\alpha_k (\lambda')^{\sigma_k}$ instead of α_k in the demand system generates the same expenditures by industry.

²²By multiplying σ_k by a common multiplier σ' and replacing λ_n by $\lambda_n^{\frac{1}{\sigma'}}$, we obtain the same demand by industry and the same total expenditures (maintaining $\lambda_{USA} = 1$).

²³We use the square root of the size of each industry as weights, given that we obtain larger standard errors for smaller industries.

²⁴This fixed- θ specification imposes a strong link between income elasticities of demand and the coefficient for Φ in the estimation of Equation (24). Note that the link between price elasticities and income elasticities holds whenever preferences are separable and is called Pigou's law (see Deaton and Muellbauer (1980)).

²⁵Pagan (1984) describes biases in estimating standard errors with generated regressors.

and final demand. To document the role of errors in the first-step regression in affecting standard errors in income elasticities, we also construct standard errors by bootstrapping the second step only, neglecting the generated-regressor issue.

3.2 One-step estimation of the two special cases

To ensure the robustness of our income elasticity estimates, our benchmark estimation framework allows for any pattern of comparative advantage by using exporter-industry fixed effects in the gravity equation. Two special cases (Skill- and Theta-driven Models) assume more specific patterns of comparative advantage, which create explicit links between supply and demand characteristics, calling for a one-step estimation with additional cross-restrictions on supply and demand.

Skill-driven Model:

In the Skill-driven Model, we remove Ricardian forces of comparative advantage by imposing common productivity across sectors. We also impose common dispersion parameters $\theta_k = \theta$. In this model, comparative advantage is solely driven by differences in skilled-labor intensity across sectors, assuming that skilled and unskilled labor are the only two factors of production.

Combining expressions for individual final demand (multiplied by population L_n) and gravity with expression (12) for S_{ik} , we obtain the following specification for the Skill-driven Model:

$$\begin{aligned} \log X_{nik} = & \theta \log z_i - \sum_f \theta \beta_{kf} \log w_{fi} - \sum_{var} \theta \delta_{var,k} TC_{var,ni} \\ & - \sigma_k \cdot \log \lambda_n + \log L_n + \log \alpha_{5,k} + \frac{(\sigma_k - 1 - \theta)}{\theta} \log \Phi_{nk} + \varepsilon_{nik}^{M2} \end{aligned} \quad (27)$$

where Φ_{nk} satisfies the following constraint:

$$\Phi_{nk} = \sum_i \exp \left[\theta \log z_i - \sum_f \theta \beta_{kf} \log w_{fi} - \sum_{var} \theta \delta_{var,k} TC_{var,ni} \right] \quad (28)$$

We simultaneously estimate demand-side parameters (σ_k , λ_n and $\alpha_{5,k}$) and supply-side parameters (factor prices w_{fk} , TFP z_i and trade cost elasticity $\delta_{var,k}$).²⁶ Trade flows are regressed using Poisson PML constrained by Equations (27) and (28), and imposing the sum of fitted expenditures to equal observed income e_n (we do not however impose any restrictions on the

²⁶The GTAP data do not provide information on factor costs. As an external validity check, we verified that countries with a higher schooling years average (Barro and Lee, 2012 update) are associated with a comparative advantage in skill-intensive industries (i.e. larger exporter fixed effects in skill-intensive industries).

trade balance). Observed variables include trade flows X_{nik} , population L_n , income e_n , trade cost variables $TC_{var,ni}$ (without exporter-specific border effects) and skilled and unskilled labor intensities β_{Hk} and β_{Lk} (normalized such that $\beta_{Hk} + \beta_{Lk}$ sum to one).

Theta-driven Model:

In the Theta-driven Model, differences in the dispersion parameter θ_k across sectors are identified by exploiting restrictions on the supply side (patterns of comparative advantage and trade costs) and the demand side (the coefficients on Φ). Following Fieler (2011), we impose the differences in trade costs elasticities to be driven by differences in θ_k and assume that δ_{var} is constant across sectors. Combining the expression for individual final demand and gravity with expression (13) for S_{ik} , we obtain our Theta-driven Model specification:

$$\begin{aligned} \log X_{nik} = & \log T_i - \theta_k \log w_i - \sum_{var} \theta_k \delta_{var} TC_{var,ni} \\ & - \sigma_k \cdot \log \lambda_n + \log L_n + \log \alpha_{5,k} + \frac{(\sigma_k - 1 - \theta_k)}{\theta_k} \log \hat{\Phi}_{nk} + \varepsilon_{nik}^{M3} \end{aligned} \quad (29)$$

with the constraint: $\Phi_{nk} = \sum_i \exp [\log T_i - \theta_k \log w_i - \sum_{var} \theta_k \delta_{var} TC_{var,ni}]$.

Observed variables include trade flows X_{nik} , population L_n , income e_n and trade costs variables $TC_{var,ni}$.²⁷ As in the Skill-driven Model and the benchmark estimations, we also impose the sum of fitted expenditure to equal observed income e_n . We estimate the following parameters: θ_k , σ_k , the Lagrangian multiplier λ_n , wage w_i ,²⁸ trade costs elasticities δ_{var} and industry fixed effects $\alpha_{5,k}$.

3.3 Data

Our empirical analysis is almost entirely based on the Global Trade Analysis Project (GTAP) version 7 dataset (Narayanan and Walmsley, 2008). GTAP contains consistent and reconciled production, consumption, endowment, trade data and input-output tables for 57 sectors of the economy, 5 production factors, and 94 countries in 2004. The set of sectors covers both manufacturing and services and the set of countries covers a wide range of per-capita income levels. The list of sectors can be found in Table 2.

To estimate gravity Equations (22) by industry, we use gross bilateral trade flows from GTAP measured including import tariffs, export subsidies and transport costs (c.i.f.). Demand systems are estimated over all 94 available countries using final demand values based on the

²⁷As for the Skill-driven Model, we do not include exporter-specific border effects.

²⁸Wages w_i are taken as free parameters but we obtain similar results using GDP per capita instead.

aggregation of private and public expenditures.²⁹ Some sectors in GTAP are used primarily as intermediates and correspond to extremely low consumption shares of final demand. 6 sectors for which less than 10% of output goes to final demand (coal, oil, gas, ferrous metals, metals n.e.c. and minerals n.e.c.) are assumed to be used exclusively as intermediates and are dropped from the final demand estimations. We also drop “dwellings” from our analysis and are left with 50 sectors.

Factor usage data by sector are directly available in GTAP and cover capital, skilled and unskilled labor, land and other natural resources. There are, however, some limitations concerning the skill decomposition of labor: while the GTAP dataset provides skilled vs. unskilled labor usage for all countries, part of this information is extrapolated from a subset of European countries and 6 non-European countries (US, Canada, Australia, Japan, Taiwan and South Korea).³⁰ Also, skilled labor is defined on an occupational basis for some of these countries (e.g. US). In most of our analysis, we measure factor intensities by the weighted average factor intensities across all countries, but our results carry on if our factor intensity measures are solely based on the subset of countries mentioned above, as shown in Appendix E.

Finally, bilateral variables on physical distance, common language, access to sea, colonial link and contiguity are obtained from CEPII (www.cepii.fr).³¹ Dummies for regional trade agreement and common currency are from de Sousa (2012).

3.4 Demand system estimation results

We focus here on the results from the 2-step estimation of the general model. Summary statistics for the two special cases (Skill- and Theta-driven Models) can be found in Appendix A. Results from the gravity equation (step 1) are standard and also presented in detail in Appendix A. In brief, there is significant variation in the distance and border effect coefficients across industries. As usually found in the gravity equation literature, the coefficient for distance is on average close to -1. Coefficients for other trade cost proxies are significant for most industries. The border effect coefficients are large, and allowing them to vary across exporters improves the model’s fit without substantially affecting the coefficients for traditional trade costs variables such as distance. As in Waugh (2010), these border effects are found to be negatively correlated with exporter per-capita income.

We now focus on the final demand estimation (step 2), Equation (24). Summary statistics

²⁹We use trade in final goods computed from GTAP using the proportionality assumption to estimate the Theta- and Skill-driven models (eqs. 27 and 29).

³⁰See: <https://www.gtap.agecon.purdue.edu/resources/download/4183.pdf>

³¹Distance between two countries is measured as the average distance between the 25 largest cities in each country weighted by population. Similarly, internal distance within a country is measured as the weighted average of distance across each combination of city pairs. See Mayer and Zignago (2011).

Table 1: NLLS estimation of final demand: regression statistics

Specification:	(1) Benchmark (Asym. TC)	(2) Symmetric trade costs	(3) $\theta = 4$	(4) No budget constraint	(5) $\Phi = 0$
Correlation σ_k with benchmark specification	1	0.916	0.913	0.998	0.946
Weighted av. coeff on Φ_{nk}	0.341	0.510	0.368	0.306	/
Correlation log λ_n with log per capita income	-0.992	-0.982	-0.979	-0.984	-0.999
Correlation θ_k with σ_k	0.110	0.201	/	0.167	/
θ_k 75th/25th pctile ratio	2.408	1.912	/	2.412	/
F-stat $\sigma_k = \sigma$	12.58	8.85	4.62	14.70	15.92
R2	0.784	0.785	0.775	0.791	0.750
Partial R2	0.279	0.281	0.219	0.316	0.150
AIC	3.025	3.023	3.047	2.965	3.169
BIC	3.360	3.358	3.314	3.300	3.436
Parameters	244	244	194	244	194
Observations	4700	4700	4700	4700	4700

Notes: Constrained NLLS regressions: step 2 of the estimation procedure described in the text; weighted by industry size (world expenditure by industry); “Partial R2” computed as $1 - \frac{SSE}{SSE^{homoth}}$; AIC (Akaike Information Criterion) computed as $\ln(\frac{SSE}{n}) + \frac{2k}{n}$ and BIC (Bayesian Information criterion) as $\ln(\frac{SSE}{n}) + \frac{k \ln(n)}{n}$, where n is the number of observations and k is the number of parameters.

are reported in Table 1. Column (1) corresponds to our benchmark specification; column (2) is identical to column (1) except that the border effects used to estimate the Φ 's are not allowed to vary across exporters (symmetric trade costs); column (3) is identical to column (1) except that θ_k is imposed to equal 4 in each sector; column (4) drops the constraint that fitted expenditures add up to observed total expenditures, and column (5) estimates demand without controlling for cross-country price differences, which is equivalent to imposing $\Phi = 0$.

In all cases, a large part of the variability in the dependent variable x_{nk} is captured by industry fixed effects, which leads to very high measures of fit (weighted R2). In order to better illustrate the contributions of non-homotheticity and price differences in explaining demand patterns, we also propose an alternative metric (Partial R2) which measures the increase in fit relative to a model with homothetic preferences and no trade costs (i.e. imposing common $\sigma_k = \sigma$ and $\Phi = 0$). This reference point also corresponds to regressing the log of expenditures on country and sector fixed effects. The partial R2 in column (1) shows that our benchmark specification captures 28% of the variability left unexplained by homothetic preferences without trade costs. In comparison, homothetic preferences with asymmetric trade costs yield a partial R2 of 0.18. In column (5), the specification with no trade costs ($\Phi = 0$) shows that non-

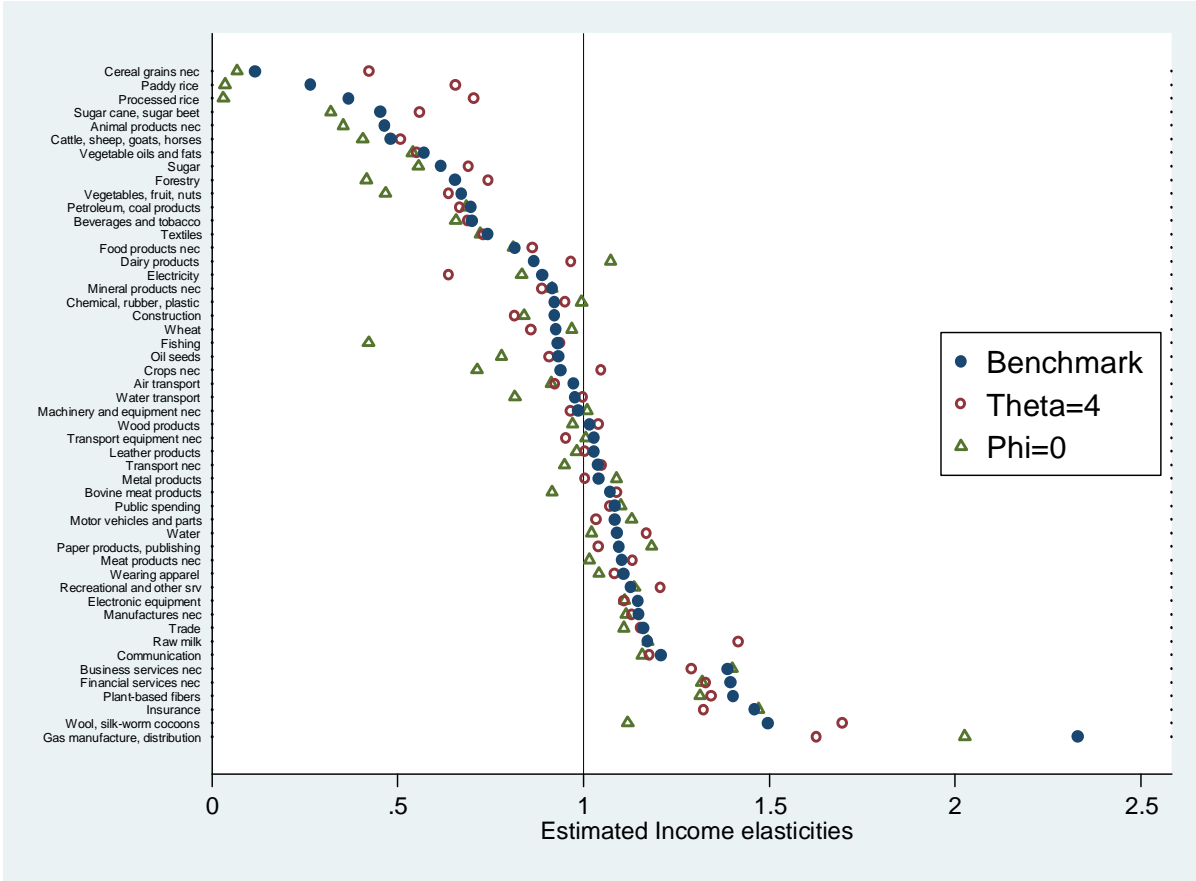


Figure 1: Income elasticity estimates across specifications

homotheticity alone captures 15% of the variability left unexplained by homothetic preferences without trade costs.

The contribution of non-homotheticity to the fit of demand patterns is statistically significant: the F-stats associated with imposing common σ_k 's across industries (sixth row of Table 1) show that homotheticity is clearly rejected in all specifications (all P-values < 0.001). Similarly, the inclusion of Φ_{nk} significantly improves this fit. In the specifications of columns (1) to (4), the coefficients associated with $\hat{\Phi}_{nk}$ are found to be jointly significant (all P-values < 0.001). Both the Akaike (AIC) and Bayesian (BIC) information criteria favor the specification that does not impose individual expenditures to equal observed income. According to both criterion, the specifications which allow for non-homothetic preferences and control for price differences (1-3) are favored to the specification that imposes no prices differences (5) as well as the specifications (not shown) with homothetic preference (with or without trade costs).³²

³²The values for AIC and BIC under homothetic preferences are 3.310 and 3.508 without controlling for prices

Table 2: Estimated income elasticity by sector

GTAP code	Sector name	Income elast.	Std error	Skill intensity
gro	Cereal grains nec	0.110*	0.133	0.135
pdr	Paddy rice	0.254*	0.199	0.061
pcr	Processed rice	0.352*	0.113	0.130
c_b	Sugar cane, sugar beet	0.433*	0.233	0.091
oap	Animal products nec	0.444*	0.098	0.132
ctl	Bovine cattle, sheep and goats, horses	0.458*	0.137	0.164
vol	Vegetable oils and fats	0.545*	0.063	0.217
sgr	Sugar	0.588*	0.085	0.221
frs	Forestry	0.623*	0.121	0.118
v_f	Vegetables, fruit, nuts	0.640*	0.136	0.095
p_c	Petroleum, coal products	0.664*	0.052	0.313
b_t	Beverages and tobacco products	0.667*	0.079	0.297
tex	Textiles	0.707*	0.064	0.231
ofd	Food products nec	0.777*	0.063	0.268
mil	Dairy products	0.826*	0.077	0.248
ely	Electricity	0.848*	0.073	0.372
mmm	Mineral products nec	0.874	0.097	0.281
crp	Chemical, rubber, plastic products	0.880	0.067	0.356
cns	Construction	0.880	0.061	0.294
wht	Wheat	0.883	0.202	0.117
fsh	Fishing	0.886	0.139	0.124
osd	Oil seeds	0.889	0.194	0.119
ocr	Crops nec	0.893	0.144	0.115
atp	Air transport	0.929	0.070	0.313
wtp	Water transport	0.932	0.100	0.299
ome	Machinery and equipment nec	0.938	0.066	0.372
lum	Wood products	0.970	0.103	0.248
otn	Transport equipment nec	0.981	0.076	0.343
lea	Leather products	0.981	0.066	0.212
otp	Transport nec	0.990	0.074	0.296
fmp	Metal products	0.992	0.077	0.297
cmt	Bovine meat products	1.023	0.078	0.238
osg	Public Administration and services	1.033	0.049	0.503
mvh	Motor vehicles and parts	1.034	0.066	0.341
wtr	Water	1.039	0.087	0.378
ppp	Paper products, publishing	1.044	0.093	0.340
omt	Meat products nec	1.052	0.096	0.233
wap	Wearing apparel	1.057	0.069	0.247
ros	Recreational and other services	1.075	0.067	0.475
ele	Electronic equipment	1.094	0.070	0.358
omf	Manufactures nec	1.095	0.065	0.279
trd	Trade	1.106	0.070	0.308
rmk	Raw milk	1.118	0.145	0.152
cmn	Communication	1.152*	0.078	0.485
obs	Business services nec	1.324*	0.059	0.504
ofi	Financial services nec	1.331*	0.090	0.546
pfb	Plant-based fibers	1.339*	0.193	0.167
isr	Insurance	1.392*	0.104	0.533
wol	Wool, silk-worm cocoons	1.426*	0.177	0.089
gdt	Gas manufacture, distribution	2.221*	0.260	0.362

Notes: Estimates based on the benchmark specification; income elasticities evaluated using median country expenditure shares; bootstrapped standard errors (500 draws); * denotes 5% significance (difference from unity); skill intensity based on total requirements.

The estimated σ_k can be used to compute income elasticities $\hat{\eta}_{nk}$ according to Equation (26). Table 2 displays estimates from the benchmark model computed using fitted median-income-country expenditure shares as weights.³³ Estimates range from 0.110 for cereal grains to 2.221 for gas manufacture and distribution with a clear dominance of agricultural sectors at the low end and service sectors at the high end. Half of the estimates are significantly different from 1 (at 95 %). Standard errors are on average equal to 0.102 when both estimation steps are run for each bootstrap. Only accounting for errors in the second step, i.e. assuming Φ_{nk} to be an error-free variable, yields an average standard error of 0.094. This small difference suggests that measurement errors stemming from the first step are small. A third alternative is to construct bootstrap by resampling countries but not sectors.³⁴ This method again yields very similar standard errors.

The distribution of estimated income elasticities is quite similar across specifications (see Figure 1). In particular, we find that the choice of θ_k does not substantially affect estimates of σ_k . As shown in Table 1, the correlation between the estimated σ_k in other specifications and those of the benchmark specification is always above 85%. This is also the correlation between income elasticities among specifications since income elasticities are proportional to σ_k . Sectors where income elasticities vary the most across specifications are actually the smallest ones (such as wool), and weighing this correlation by final demand yields larger correlation estimates in all cases.

For robustness, our estimated income elasticities are compared with estimates based on AIDS and LES, two more standard demand systems, and are found to be well correlated (Appendix E). In addition, we propose a reduced-form approximation of our benchmark equation (Appendix B). Since the Lagrangian multiplier λ_n is highly negatively correlated with per capita income (in log), we can approximate income elasticities using coefficients on log per-capita income instead of the log of the Lagrangian and find similar estimates.

In the benchmark specification, we can also use our estimates of σ_k to examine the differences in θ_k across sectors implied by the estimated coefficient on Φ_{nk} . Doing so, we find a positive but not statistically significant correlation between θ_k and σ_k (fourth line of Table 1).

Alternatively, we can also use the Skill- and Theta-driven Models to estimate σ_k and θ_k across sectors. While we leave for Appendix A the summary statistics of the estimation of each special case, note that we find the correlation of the estimated σ_k with our benchmark

and 3.136 and 3.403 if controlling for them. AIC yields a larger gap between homothetic and non-homothetic preferences as it puts a smaller penalty on models with more degrees of freedom.

³³With CRIE preferences, the ratio of income elasticities between two sectors does not depend on the choice of the reference country.

³⁴This “block-bootstrap” approach accounts for clusters if errors are correlated across industries for each importer.

estimates to be 60% for the Skill-driven Model and 61% for the Theta-driven Model. Moreover, the correlation between σ_k and θ_k in the Theta-driven Model is negative at -0.16 but not significant at the 10% level.

3.5 Correlation with factor intensities

We now investigate the relationship between income elasticities and factor intensities across sectors. Although the implications of such a relationship will be best illustrated in Section 4, we first demonstrate its significance through simple correlations. Table 3 reports correlation coefficients between skill intensity and income elasticity, or, in columns 2 and 4, the beta coefficients associated with each intensity parameter in regressions of income elasticity on several factor intensities. It displays standard errors constructed by resampling importers and sectors in all steps of the estimation: the two steps required to estimate income elasticities as well as the correlation with factor intensities.³⁵

Our measures of factor intensity correspond to the ratio of skilled labor, capital or natural resource (including land) to total labor input. They are computed including the factor usage embedded in the intermediate sectors used in each sector’s production, based on data pooled across all countries for greater precision. Appendix E shows that our results are robust to different measures of factor intensities. Table 3 reports estimates resulting from CRIE preferences, while alternative demand systems are examined in Appendix D. The correlation with skill intensity is also illustrated in Figure 2.

We find that skill intensity is positively and significantly correlated with income elasticity, natural resources intensity is weakly negatively correlated, and capital intensity exhibits a weakly positive correlation. As expected, the correlation with skill intensity diminishes if we account for trade costs and control for differences in price indexes. This can be seen by comparing column (1) versus (6) in Table 3. This correlation remains however particularly large and around 50% in most specifications.

This correlation is not driven by sectors that contribute little to final demand: it is even stronger when we weight observations by world final consumption in each industry (not shown). Part of this large correlation can be explained by the composition of consumption into services

³⁵We compare them to bootstrapped standard errors resulting from taking the Φ ’s as perfectly measured and find similar results: for example, the estimate in column 1 is 0.121 instead of 0.120. Alternatively, we have computed standard errors on the correlation coefficient using a feasible generalized least squares (FGLS) regression in which the bootstrapped standard errors from the NLLS estimations of income elasticities are used to construct weights (see Lewis and Linzer (2005)). These lead to standard error estimates (0.113) which are very close to both those resulting from the full bootstrapping (0.120) and to those resulting from a simple robust OLS regression of income elasticity on skill intensity (0.123). The similarity between estimates suggests that the bias caused by the use of generated variables is small.

Table 3: Correlation between income elasticity and skill intensity

Dep. var.:	Income elasticity							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Specification	Benchmark (Asym TC)	Benchmark (Asym TC)	Symmetric trade costs	Symmetric trade costs	$\theta = 4$	$\Phi = 0$	Excl. services	Excl. services
Skill intensity	0.523 [0.120]**	0.531 [0.158]**	0.496 [0.118]**	0.465 [0.163]**	0.513 [0.161]**	0.651 [0.088]**	0.361 [0.181]*	0.490 [0.231]**
Capital int.		0.068 [0.268]		0.062 [0.233]				-0.385 [0.356]
Nat. rces. int.		-0.005 [0.180]		-0.203 [0.145]				0.213 [0.290]
Obs. (sectors)	50	50	50	50	50	50	42	42

Notes: Dependent variable: income elasticity by sector evaluated at median-country income; “Nat. rces. int.”: natural resources intensity; beta coefficients; bootstrap standard errors in brackets (500 draws); * significant at 5%; ** significant at 1%.

vs. manufacturing industries, with the former being generally associated with a larger income elasticity. However, columns (7) and (8) of Table 3 show that the correlation remains high in all specifications even after excluding service industries. In that case, the correlation is significant at 5% (column 7) and 1% (column 8) depending on whether we control for capital intensity and natural resources intensity. We also test the correlation between skill intensity and income elasticity using estimates from the two special cases (Skill- and Theta-driven Models) where we impose additional constraints on the supply side. We find even stronger correlations: 65.3% for the Skill-driven Model and 71.6% for the Theta-driven Model.³⁶

It is interesting to note that capital intensity would otherwise be positively correlated with income elasticity, as found by Reimer and Hertel (2010), but this correlation is not as large as for skill intensity (less than 10% in most specifications) and not robust to controlling for skill intensity as shown in columns (2) and (4) of Table 3.

These results imply a large correlation between per capita income and final demand in skill-intensive sectors. We emphasize the demand side. One may be worried, however, that these results are driven by differences in skill endowment across countries rather than differences in per capita income. In the GTAP data, the fraction of payment to skilled labor is indeed correlated at 88% with log per capita income.³⁷ In order to check the robustness of our results with respect to differences in education, we re-estimated income elasticities for subsets of countries with smaller

³⁶Note that the Theta-driven model assumes that there is only one factor of production. It shows however that alternative estimates of income elasticities are still correlated with observed skill intensity across sectors.

³⁷Similarly, the correlation between average schooling years and per capita income is very high: 74.7% using data from Barro and Lee (2012 update).

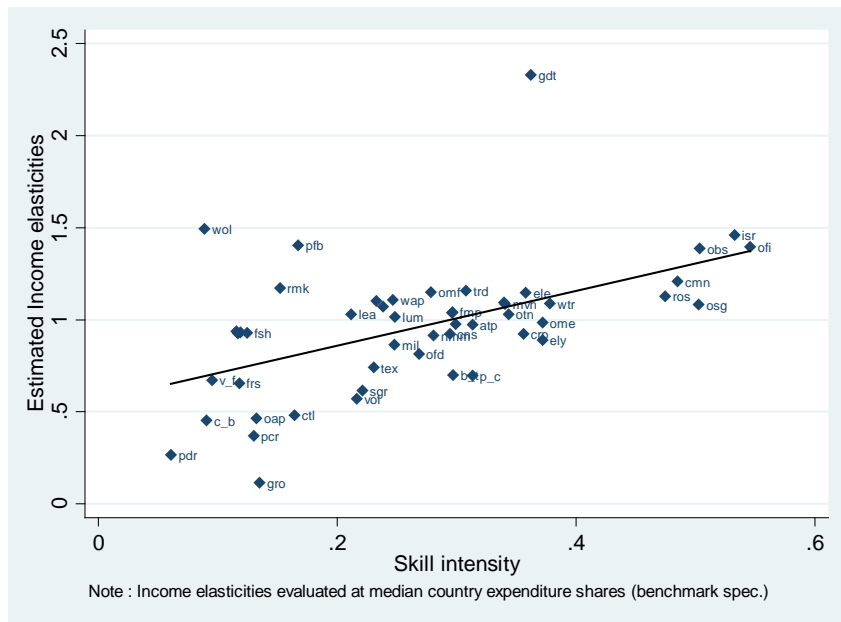


Figure 2: Income elasticity and skill intensity correlation.

variations in skilled labor endowment (and still large variations in per capita income). If we restrict the set of countries to those within the inter-quartile range in skilled-labor endowments, the correlation between estimated income elasticities and skill intensity remains very high for the main specifications (above 40%) while the correlation between per capita income and education is sensibly lower. A more extreme exercise is to select specific groups of countries where the correlation between income and education becomes zero by construction. In these cases we find again very large correlations between skill intensity and (re-estimated) income elasticity, showing that our main results are not driven by differences in education across countries.

4 Implications for trade puzzles

In the following section, we use our estimates to illustrate the implications of non-homothetic preferences for: i) the correlation between consumption and production, ii) the “missing trade” puzzle in factor content, iii) the volume of trade flows between rich and poor countries and iv) aggregate trade-to-GDP ratios. We investigate these topics using our benchmark specification (incorporating asymmetric trade costs) and an identical model with homothetic preferences as a point of reference. iii) and iv) are also investigated using the two special cases (Skill- and Theta-driven Models) in order to distinguish the roles of differences in factor intensities and differences in dispersions of productivity.

4.1 Consumption patterns and missing trade

The correlation between skill intensity and income elasticity implies that the factor content of consumption varies systematically with income. In Figure 3, we plot per capita income (in log) against a measure of the relative skilled-labor content of consumption:

$$\frac{\sum_k \beta_{Hk} \widehat{D}_{nk}}{\sum_k (\beta_{Lk} + \beta_{Hk}) \widehat{D}_{nk}} \quad (30)$$

where β_{Hk} and β_{Lk} are defined as the skilled-labor and unskilled labor intensity of production (using average of total requirements across all countries). We define final demand by using either actual consumption or fitted consumption \widehat{D}_{nk} with different assumptions. With homothetic preferences and no trade costs, expression (30) would be the same for all countries. Trade costs (including asymmetric trade costs) already explain part of the variations in the factor content of consumption: rich countries tend to spend more on skilled-labor intensive industries, even if preferences are homothetic, because goods from these industries are relatively cheaper in these countries. However, we can see in Figure 3 that an even better fit is obtained when non-homothetic preferences are allowed on top of trade costs.




Figure 3: Skilled-labor content of consumption and per capita income

As shown in Figure 4, rich countries also tend to specialize in skill-intensive sectors. This generates a correlation between relative specializations in consumption and production, which we illustrate by looking at the relationship between $\frac{V_{nk}}{s_n V_k}$ and $\frac{D_{nk}}{s_n Y_k}$ in the first row of Table 4.




Figure 4: Skilled-labor content of production and exports against per capita income

The term $\frac{V_{nk}}{s_n V_k}$ reflects actual value-added in sector k of country n relative to world total value-added in sector k multiplied by country n 's share of world expenditures. The second term $\frac{D_{nk}}{s_n Y_k}$ corresponds to the relative specialization in consumption which is computed using fitted final demand \widehat{D}_{nk} from our second-stage estimates including trade costs (both symmetric and asymmetric) in columns (1) to (3) and using observed consumption D_{nk} in column (4). Note that if preferences were homothetic and if there were no trade costs, as is assumed in standard Heckscher-Ohlin models, the correlation would be null as consumption patterns would be the same across all countries and $\frac{D_{nk}}{s_n Y_k} = 1$. In columns (1) and (2), we find a positive correlation between consumption and production even if preferences are assumed to be homothetic. The estimated correlation across countries and industries is 34% and 37% (with and without asymmetric border effects) and significantly positive at 1%. Column (3) shows that this correlation increases to 49% if we account for both non-homothetic preferences and trade costs. This value is closer to the 62% correlation observed in the data (column 4).

Missing trade in factor content: slope and variance tests

A positive correlation between income elasticity and skill intensity not only generates a correlation between supply and demand, but also generates a smaller factor content trade compared to the homothetic case.

As described in Section 2.4.3, the predicted factor content of trade (PFCT) can be ex-

Table 4: Patterns of supply, demand and the factor content of trade

	(1)	(2)	(3)	(4)	Dimension
Preferences:	Homothetic		Non-Homoth.	data	
Asymmetric trade costs:	No	Yes	No		
1- Correlation between supply and demand $\frac{V_{nk}}{s_n V_k}$ and $\frac{D_{nk}}{s_n D_k}$	0.365	0.344	0.485	0.619	n x k
2- Correlation between F_{nf}^{HOV} and Consumption bias F_{nf}^{CB}	0.582	0.664	0.773	0.793	n x f
3- Skilled/unskilled labor only	0.473	0.383	0.573	0.648	n x 2
Corrected HOV slope test:					
4- Assuming common β_{kf}	0.708	0.767	0.887	1	n x f
5- Using global input-output matrix	0.476	0.537	0.680	1	n x f
Missing trade: $\frac{Var(F_{fn}^{meas})}{Var(F_{fn}^{pred})}$					
6- Assuming common β_{kf}	0.670	0.747	0.890	1	n x f
7- Using global input-output matrix	0.450	0.452	0.681	1	n x f

pressed as the difference between standard Heckscher-Ohlin-Vanek PFCT, denoted F_{nf}^{HOV} , and a consumption bias term denoted F_{nf}^{CB} (see equation 17). Assuming constant requirements coefficients β_{kf} across countries, we impute F_{nf}^{HOV} using production data and F_{nf}^{CB} using either fitted final demand (columns 1 to 3) or actual consumption (column 4).³⁸ This correlation between F_{nf}^{HOV} and F_{nf}^{CB} would be null if preferences were homothetic and there were no trade costs ($F_{nf}^{CB} = 0$). The second row of Table 4 shows that trade costs would generate a large correlation between the factor content of consumption and production even if preferences were assumed to be homothetic (column 1). This correlation is 58% across countries and factors. This is consistent with Davis and Weinstein (2001) who also attribute an important part of the missing trade puzzle to trade costs. Accounting for asymmetric trade costs further reduces trade between rich and poor countries and generates a larger correlation between the factor content of consumption and production as shown in column (2). In column (3), we find however that allowing for non-homotheticity further increases the correlation to 77%, which is even closer to the large correlation found using observed final demand (79%).

³⁸For this correlation, as well as the slope and variance tests, all observations are scaled by $(s_n \sum_i w_{if} V_{if})^{1/2}$ to adjust for heteroskedasticity. Note also that all variables are in value terms (e.g. wages instead of number of workers) which mitigates cross-country differences related to differences in factor prices.

These correlations have important implications for the “missing trade” puzzle. In rows 4 to 7, we examine the “slope test” and the “variance test” traditionally conducted to test the Heckscher-Ohlin-Vanek model and amended versions pioneered by Trefler (1995). The slope test is simply the coefficient of the regression of the measured factor content of trade on the predicted factor content. The variance test is the ratio of the variances of the measured and predicted factor contents of trade. The latter best reflects the “missing trade” puzzle: previous results have found small ratios (Trefler 1995). Both tests would exhibit a coefficient equal to one if the predicted and measured factor contents were equal.

We construct the predicted and measured factor contents of trade in two ways. In rows 4 and 6 of Table 4, we follow the same strategy as in row 2 by assuming constant factor requirement coefficients across countries. In rows 5 and 7, the factor content is computed by accounting for trade in intermediate goods and differences in factor requirement across countries. Following the method developed by Trefler and Zhu (2010), we construct a matrix of direct and indirect factor requirements by taking into account factors embodied in traded intermediate goods. Data on domestic and imported input requirements at the country level are provided in the GTAP database. In Appendix G, we provide additional details on the measurement of the factor content of trade as well as a generalization of the Trefler and Zhu (2010) framework in which the “consumption similarity” condition is violated by trade costs or non-homothetic preferences. In this case, the measured factor content of trade corresponds to the “actual” factor content (as defined by Trefler and Zhu 2010).

In row 4 of Table 4, assuming common input-output coefficients, we find that non-homothetic preferences push the slope coefficient from 0.708 in column (1) with homothetic preferences (0.767 with asymmetric trade costs) to 0.887 in column (3), bringing it 61% closer to unity (48% compared to asymmetric trade costs). In row 5, accounting for intermediate goods and country-specific input-output coefficients, non-homothetic preferences increase the slope coefficient from 0.476 in column (1) (0.535 with asymmetric trade costs) to 0.680 in column (3). This corresponds to a 39% reduction in the gap towards a unit coefficient (31% reduction from asymmetric trade costs).

Rows 6 and 7 of Table 4 display results from the “variance test” . When we assume common input requirement coefficients, allowing for non-homotheticity improves the ratio from 0.670 in column (1) (0.747 with asymmetric trade costs in column 2) to 0.890 in column (3). This actually corresponds to a 75% decrease in the excess variance predicted by the model compared to the variance in measured factor content of trade. When we account for intermediate goods trade and country-specific requirements in inputs, the variance ratio increases from 0.450 with homothetic preferences (0.453 with asymmetric trade costs) to 0.681 with non-homothetic preferences, as shown in row 7 of Table 4. This improvement corresponds to a 60% decrease in

the excess variance of the predicted factor content of trade relative to the measured factor content of trade. Together, results from the slope and variance tests suggest that non-homothetic preferences can explain about half of the missing trade puzzle.³⁹

4.2 Trade volumes and trade partners

Results from the previous section shed light on the role of non-homothetic preferences in explaining the net factor content of trade. In particular, our results are related to the industry composition of demand and production. Given that a large fraction of trade is intra-sectoral, it is important to ask whether non-homotheticity can also play a role (quantitatively) in explaining patterns of gross trade volumes.

To examine trade volumes between countries, we use estimates from both steps of our benchmark estimation. We reconstruct trade shares using estimates from the full gravity equation with exporter-specific border effects (asymmetric trade costs). According to Equation (6), we compute the fraction of goods k consumed in n that are purchased from i as $\hat{\pi}_{nik} = \frac{\hat{S}_{ik}(\hat{d}_{nik})^{-\theta_k}}{\hat{\Phi}_{nk}}$. We reconstruct total demand from fitted final demand \hat{X}_{nk} and the fitted intermediate good demand implied by input-output linkages described in Equation (5), excluding service industries. Fitted trade flows are then assembled using:

$$\hat{X}_{nik} = \hat{\pi}_{nik} \hat{X}_{nk} \tag{31}$$

We compare trade implied by non-homothetic and homothetic preferences.⁴⁰

North-South trade volumes

Can non-homothetic preferences explain why the volumes of North-South trade in comparison to North-North trade are small?⁴¹ As argued in Section 2.4.1, non-homotheticity can potentially explain differences in import penetration across markets depending on the importer’s income and the exporter’s structure of comparative advantage. Since income elasticity is correlated with skill intensity, richer countries tend to consume relatively more skill-intensive goods, as

³⁹For illustrations on the missing trade puzzle, we chose to include service industries as in most of the literature including Trefler (1995), Davis and Weinstein (2001) and Trefler and Zhu (2010). If we exclude services, our results are slightly weakened but still exhibit a strong role for non-homothetic preferences.

⁴⁰For both cases (non-homothetic and homothetic preferences), the fitted trade shares allow for asymmetric trade costs as in Waugh (2010) to distinguish their effect from those implied by non-homothetic preferences. In this subsection, we exclude services as most papers in the literature (including Fielser 2011) to avoid concerns about mismeasurement of bilateral trade flows in service industries. Note, also, that we do not impose trade balance in our estimation. The ratio of trade deficit to GDP implied by the fitted trade flows is however highly correlated across countries with the ratio found in the data.

⁴¹See Fielser (2011), Waugh (2010) among others.

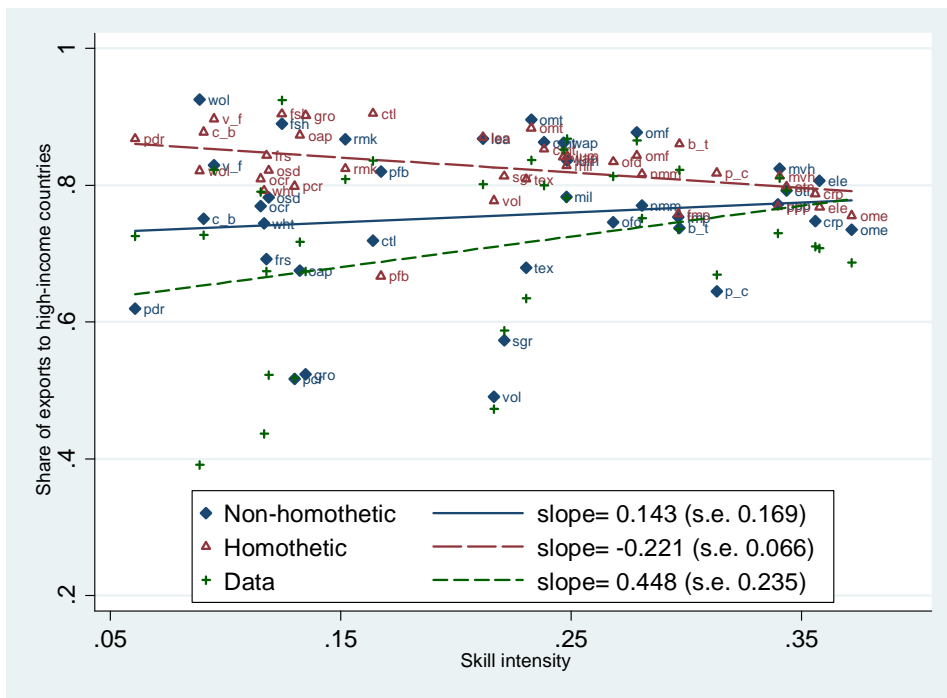


Figure 5: Share of imports by rich partners across sectors depending on skill intensity

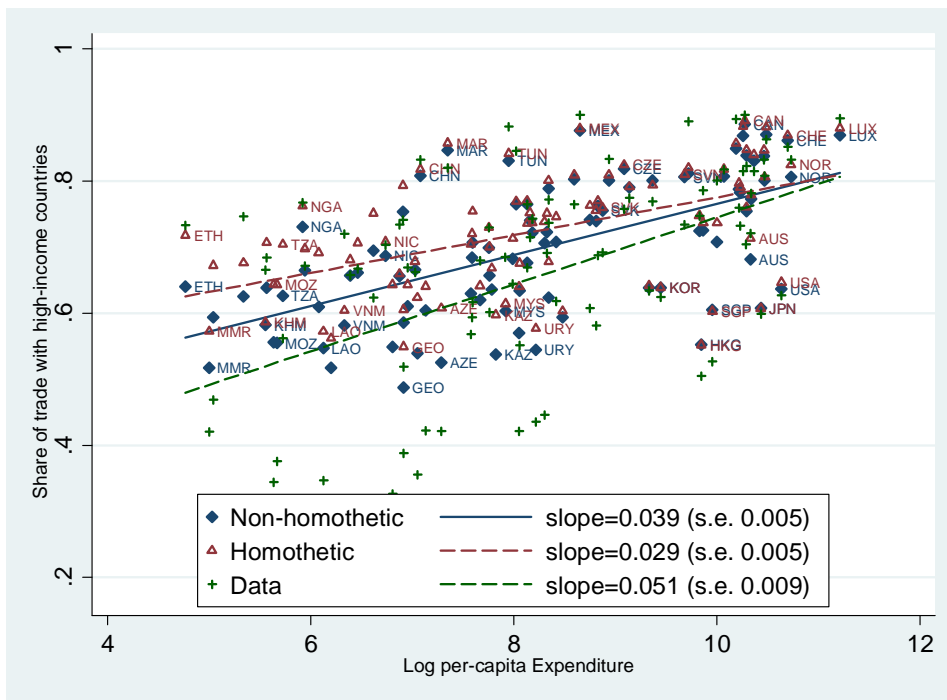


Figure 6: Share of trade with rich partners across countries (imports and exports)

shown in Figure 3. We find a similar result for the patterns of imports: high-income countries tend to import relatively more skill-intensive goods, especially compared to the case with homothetic preferences. Figure 5 shows that rich countries tend to be the main destination for the international trade of skill-intensive goods even if they have a comparative advantage at producing these goods. This implies that the demand channel offsets the effect of competition.

Figure 6 plots each country’s share of trade with high-income partners (defined as having annual per capita income above \$10K). As we can see, homothetic preferences with trade costs can already generate a positive correlation with per-capita income since richer countries tend to be relatively close to each other and therefore more likely to trade together. As expected, however, non-homothetic preferences further magnify this correlation and improve the fit with the data (slope coefficient of 0.039 for non-homothetic preferences compared to 0.029 for homothetic preferences and 0.051 in the data). In particular, we can observe substantial differences in predicted shares for the poorest countries. For the lowest-income countries, predicted trade shares with rich countries are 10% lower with non-homothetic preferences than with homothetic preferences.




Figure 7: Share of trade with rich partners: Skill-driven vs. Theta-driven Models

While we find a significant role for non-homothetic preferences in determining trade patterns, the above estimates are based on our benchmark model which, by allowing for any pattern of comparative advantage, is very flexible on the supply side. To further pinpoint the

role of the skill-intensity-to-income-elasticity correlation, we compare fitted trade flows from the Skill-driven and Theta-driven Models. Figure 7 shows that the Skill-driven Model generates a stronger relationship between per capita income and the share of trade with rich countries than the Theta-driven Model (with a slope coefficient of 0.044 compared to 0.034 for the latter). These results indicate that non-homotheticity on the demand side interacts more with patterns of comparative advantage stemming from differences in skilled vs. unskilled labor rather than from differences in the dispersion coefficients θ_k .

The home bias puzzle

Similarly, we use our estimates from the first and second stage regressions to construct fitted trade-to-GDP ratios and plot them against per capita income (in log) in Figure 8. As the model allows for asymmetric trade costs which decrease with income as in Waugh (2010), it already predicts an increasing relationship with homothetic preferences.⁴² Despite this, the relationship between income and openness to trade is even stronger when allowing for non-homotheticity. In this case, fitted trade flows well replicate the relationship found in the data. In particular, the trade-to-GDP ratio is smaller for low-income countries when we allow for non-homotheticity in preferences. Interestingly, this ratio is larger for rich countries.




Figure 8: Openness across countries $(\text{Export}/\text{Output} + \text{Imports}/\text{Absorption})/2$

⁴²The relationship would otherwise be flat or negative without allowing for exporter-specific border effects.

In Section 2.4.2, we argue that at least two mechanisms might explain the impact of non-homotheticity on openness to trade. To examine them, we follow the decomposition proposed in Equation (16). Figure 9 plots the first and second term of expression (16) as a function of log per capita income. Each term relates to the difference in the domestic share of trade between non-homothetic and homothetic fitted trade flows. The first term illustrates the “covariance” channel: an increase in covariance between consumption patterns and comparative advantage when preferences are non-homothetic. As expected, it is positive on average and for more than half of the countries. It is also strongest for countries at both ends of the income distribution; those are the countries for which consumption patterns are most dissimilar to other countries. The second term indicates that income-elastic goods are systematically more traded, which induces richer countries to have higher openness ratios.⁴³ Figure 9 shows that this channel can better explain why poorer countries tend to have smaller trade-to-GDP ratios.

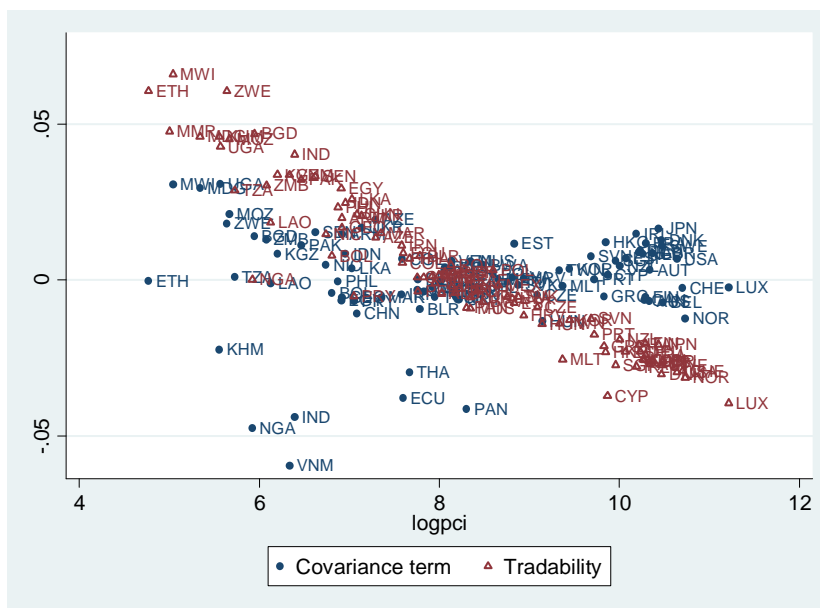


Figure 9: Decomposition of differences in openness

Can the Theta-driven Model based on Fieler (2011) better explain the trade-to-GDP ratios across countries? Interestingly, we find that the Skill-driven and Theta-driven Models yield very similar predictions. Both models (with symmetric trade costs) yield similar average levels of trade-to-GDP ratios with a flat or downward slopping relationship between per capita income and trade-to-GDP ratios. Allowing for asymmetric trade costs also does not favor one model against the other, showing that both mechanisms have similar explanatory powers.

⁴³As shown in Appendix A, this relationship is better explained by the correlation of income elasticity with the elasticity of trade costs to distance than with θ_k across sectors.

5 Summary and conclusions

We begin with the assertion that a large proportion of both theoretical and empirical research on international trade focuses on the production side of general equilibrium. The purpose of this paper is then to demonstrate that an examination of the role of demand can contribute to explaining a number of persistent puzzles long debated by trade economists. In particular, we are interested in the systematic relationship between certain characteristics of demand and characteristics of goods and services in production.

Our first task is to develop and estimate a model in which preferences are assumed to be identical across countries but non-homothetic. It allows goods to differ in their income elasticity of demand and expenditure shares to be related to per-capita income. Both economically and statistically, we find large deviations of income elasticity estimates from the unitary values implied by homothetic preferences. The next step is to relate these income elasticities of demand to factor intensities of goods in production. Here we find a strong, positive correlation (higher than 45 percent) between a good's income elasticity of demand and its skilled-labor intensity in production. The correlation is robust to the inclusion of trade costs and a number of other factors.

We then investigate the implications of non-homothetic preferences and their relationship to factor intensities. Our first results assess their contribution to the “missing trade” puzzle. We find that they can reduce the overpredicted variance in the factor content of trade by more than half. This result is driven by a supply-demand correlation which is absent under homothetic preferences: countries tend to specialize in the consumption of the same goods that they are specialized in producing.

Another set of results relate to trade patterns and the selection of trading partners. Our findings imply that high-income countries have a comparative advantage in high-income-elasticity goods and services, because these are skilled-labor intensive and because the high-income countries are skilled-labor abundant. This suggests that countries should be more likely to trade with countries of similar income level and we verify that this is the case. While this mechanism also predicts lower levels of trade-to-GDP ratios, we find its explanatory power to be smaller than another channel through which non-homotheticity affects openness to trade: the fact that income elastic goods are systematically more tradable. Results show that the two channels together are capable of closely replicating the positive relationship between per capita income and openness found in the data.

While the correlation between skill intensity and income elasticity affects trade patterns and trade volumes, we would like to point out that it may also have important implications for the skill premium (skilled worker wages relative to unskilled worker wages). In particular, it

can generate a positive effect of total factor productivity (TFP) growth on this premium. The intuition is simple. As productivity increases, people become richer and they consume more goods from income-elastic industries which are, as we show, more intensive in skilled labor. This increases the demand for skilled labor relative to unskilled labor and thus increases the relative wage of skilled workers. In the working paper version of this article (Caron et al., 2012), we used our model and estimated parameters to simulate productivity (TFP) growth and found a large effect on the skill premium, which lead us to believe that the demand-side could potentially play an important role.⁴⁴

Finally, non-homothetic preferences may also influence the direction of technological change. If income-elastic goods tend to be consumed and produced in countries that are abundant in skilled labor, the production of these goods should be associated with a larger demand for technologies compatible with skilled labor. This mechanism could partially explain the positive correlation between skill intensity and income elasticity documented in this paper.

References

- Anderson, James E. and Eric van Wincoop**, “Gravity with Gravitas: A Solution to the Border Puzzle,” *American Economic Review*, 2003, *93*(1), 170–192.
- Athey, Susan**, “Monotone Comparative Statics Under Uncertainty,” *The Quarterly Journal of Economics*, 2002, *117* (1), 187–223.
- Bergstrand, Jeffrey**, “The Heckscher-Ohlin-Samuelson Model, the Linder Hypothesis and the Determinants of Bilateral Intra-industry Trade,” *Economic Journal*, 1990, *100*(403).
- Bernasconi, Claudia**, “Income Similarity and Bilateral Trade Flows,” 2011. University of Zurich working paper.
- Caliendo, Lorenzo and Fernando Parro**, “Estimates of the Trade and Welfare Effects of NAFTA,” November 2012. NBER Working Papers 18508.
- Caron, Justin, Thibault Fally, and James R. Markusen**, “Skill Premium and Trade Puzzles: a Solution Linking Production and Preferences,” June 2012. NBER Working Papers 18131.
- Cassing, James and Shuichiro Nishioka**, “Nonhomothetic Tastes and Missing Trade in Factor Services,” 2009. Working paper.
- Chaney, Thomas**, “Distorted Gravity: The Intensive and Extensive Margins of International Trade,” *American Economic Review*, 2008, *98*(4), 1707–21.

⁴⁴The argument is not limited to the reallocation of consumption across industries: within-industry heterogeneity in product quality and income elasticity can play a similar role if the production of higher-quality products is more intensive in skilled labor.

- Costinot, Arnaud**, “An Elementary Theory of Comparative Advantage,” *Econometrica*, 2009, 77 (4), 1165–1192.
- , **Dave Donaldson**, and **Ivana Komunjer**, “What Goods Do Countries Trade? A Quantitative Exploration of Ricardo’s Ideas,” *Review of Economics Studies*, 2012, 79 (2), 581–608.
- Davis, Donald R. and David E. Weinstein**, “An Account of Global Factor Trade,” *American Economic Review*, 2001, 91, 1423–1453.
- de Sousa, Jose**, “The currency union effect on trade is decreasing over time,” *Economics Letters*, 2012, 117 (3), 917–920.
- Deaton, Angus and John Muellbauer**, “Economics and Consumer Behavior,” *Cambridge University Press*, 1980.
- Donaldson, Dave**, “Railroads of the Raj: Estimating the Impact of Transportation Infrastructure,” forthcoming.
- Eaton, Jonathan and Samuel Kortum**, “Technology, Geography, and Trade,” *Econometrica*, 2002, 79(5), 1453–1498.
- , – , and **Francis Kramarz**, “An Anatomy of International Trade: Evidence From French Firms,” *Econometrica*, 2011, 79(5), 1453–1498.
- Fajgelbaum, Pablo, Gene M. Grossman, and Elhanan Helpman**, “Income Distribution, Product Quality, and International Trade,” *Journal of Political Economy*, 2011, 119 (4), 721 – 765.
- Fally, Thibault**, “Structural Gravity and Fixed Effects,” 2012. University of Colorado Mimeo.
- , **Rodrigo Paillacar**, and **Cristina Terra**, “Economic geography and wages in Brazil: Evidence from micro-data,” *Journal of Development Economics*, 2010, 91(1), 155–168.
- Fieler, Ana Cecilia**, “Non-Homotheticity and Bilateral Trade: Evidence and a Quantitative Explanation,” *Econometrica*, 2011, 79(4), 1069–1101.
- Hallak, Juan Carlos**, “Product quality and the direction of trade,” *Journal of International Economics*, January 2006, 68 (1), 238–265.
- , “A Product-Quality View of the Linder Hypothesis,” *The Review of Economics and Statistics*, 2010, 92 (3), 453–466.
- and **Peter K. Schott**, “Estimating Cross-Country Differences in Product Quality,” *The Quarterly Journal of Economics*, 2011, 126 (1), 417–474.
- Hanoch, Giora**, “Production and Demand Models with Direct or Indirect Implicit Additivity,” *Econometrica*, 1975, 43(3), 395–419.
- Head, Keith and Thierry Mayer**, “Regional wage and employment responses to market potential in the EU,” *Regional Science and Urban Economics*, 2006, 36 (5), 573–594.

- Hung-po, Sehun Kim Chao and Alan S. Manne**, “Computation of Equilibria for Nonlinear Economies: Two Experimental Model,” *Journal of Policy Modeling*, 1982, 4(1), 23–43.
- Hunter, Linda**, “The Contribution of Non-homothetic Preferences to Trade,” *Journal of International Economics*, 1991, 30, 345–358.
- and **James R. Markusen**, *Per-Capita Income as a Determinant of Trade*, Cambridge: MIT Press,
- Khandelwal, Amit**, “The Long and Short (of) Quality Ladders,” *Review of Economic Studies*, 2010, 77 (4), 1450–1476.
- Lewis, Jeffrey B. and Drew A. Linzer**, “Estimating Regression Models in which the Dependent Variable is based on Estimates,” *Political Analysis*, 2005, 13(4), 345–364.
- Linder, Staffan B.**, *An Essay on Trade and Transformation*, Stockholm: Almqvist and Wiksell, 1961.
- Markusen, James R.**, “Explaining the Volume of Trade: An Eclectic Approach,” *American Economic Review*, 1986, 76, 1002–1011.
- , “Putting per-capita income back into trade theory,” *Journal of International Economics*, 2013, 90 (2), 255–265.
- Martinez-Zarzoso, Inmaculada and Sebastian Vollmer**, “Bilateral Trade Flows and Income-Distribution Similarity,” 2011. BE Working Papers on International Economics and Finance.
- Matsuyama, Kiminori**, “A Ricardian Model with a Continuum of Goods under Non-Homothetic Preferences: Demand Complementarities, Income Distribution, and North-South Trade,” *Journal of Political Economy*, 2000, 108, 1093–2000.
- Mayer, Thierry and Soledad Zignago**, “Notes on CEPII’s distances measures: The GeoDist database,” Working Papers 2011-25, CEPII research center Dec 2011.
- Narayanan, Badri G. and Terrie L. Walmsley**, “Global Trade, Assistance, and Production: The GTAP 7 Data Base,” 2008. Center for Global Trade Analysis, Purdue University.
- Pagan, Adrian**, “Econometric Issues in the Analysis of Regressions with Generated Regressors,” *International Economic Review*, February 1984, 25 (1), 221–47.
- Redding, Stefan James. and Anthony J. Venables**, “Economic geography and international inequality,” *Journal of International Economics*, 2004, 62(1), 53–82.
- Reimer, Jeffery J. and Thomas W. Hertel**, “Nonhomothetic Preferences and International Trade,” *Review of International Economics*, 2010, 18, 408–425.
- Simonovska, Ina**, “Income Differences and Prices of Tradeables,” 2010. NBER Working Paper No. 16233.

- **and Michael Waugh**, “The Elasticity of Trade: Estimates and Evidence,” 2010. UC Davis Working paper No. 11-2.
- Trefler, Daniel**, “The Case of Missing Trade and Other Mysteries,” *American Economic Review*, 1995, *85*, 1029–1046.
- **and Susan Zhu**, “The structure of factor content predictions,” *Journal of International Economics*, 2010, *82*, 195–207.
- Verhoogen, Eric A.**, “Trade, Quality Upgrading, and Wage Inequality in the Mexican Manufacturing Sector,” *The Quarterly Journal of Economics*, 05 2008, *123* (2), 489–530.
- Waugh, Michael**, “International Trade and Income Differences,” *American Economic Review*, 2010, *100*(5), 2093–2124.

APPENDIX

A) Gravity equation estimates

Benchmark model

Table A1 below presents the results of the gravity equation estimations (Equation 22). The first column shows the average estimated coefficient across industries while the second column shows the standard deviation of the coefficient estimate across industries. These standard deviations reflect the variations of the coefficients across industries but do not reflect measurement errors: all coefficient estimates are significantly different from zero for most industries.

Table A1: Coefficients from the gravity equation estimations

Trade cost variable:	Mean	Standard Dev.	Mean	Standard Dev.
Distance (log)	-0.886	0.698	-0.879	0.636
Contiguity	0.518	0.506	0.328	0.460
Common language	0.001	0.379	0.407	0.370
Colonial link	0.191	467	0.320	0.534
Both access to sea	0.532	0.620	0.574	0.610
RTA	0.459	0.493	0.567	0.589
Common currency	0.106	0.746	0.586	1.034
Common legal origin	0.222	0.221	0.024	0.264
Border effect	/	/	3.767	2.128
Trade costs:	Asymmetric		Symmetric	
Exporter FE	Yes		Yes	
Importer FE	Yes		Yes	
Nb. of industries	56		56	
Pseudo-R2 (incl. domestic)	0.999		0.999	
Pseudo-R2 (excl. domestic)	0.921		0.833	

Notes: Poisson regressions; dependent variable: trade flows; step 1 of the estimation procedure described in the text. The coefficients above are estimated separately for each industry. Pseudo-R2 equal the square of the correlation coefficient between fitted and observed trade flows, including or excluding domestic flows.

In the benchmark specification with asymmetric trade costs, the correlation between trade cost coefficients and income elasticity is the largest for the distance coefficient (whether we exclude services or not), and the distance coefficient is the only trade cost coefficient with a 1%-significant correlation with income elasticity. This coefficient, $\beta_{dist,k} = \theta_k \delta_{dist,k}$, is the product of the elasticity of trade to trade costs (θ_k) and the elasticity of trade costs to distance. As we show in the text, income elasticity is not significantly correlated with θ_k , which implies that income elasticity is rather correlated with $\delta_{dist,k}$. For the benchmark specification, this

correlation is 37.9% (excluding services). In other words, the effect of distance on trade costs is smaller for income-elastic goods.

Table A2: Correlations with income elasticity

Cross-sectoral correlation with income elasticity		
Correlation with distance coeff	-0.418	Asymmetric TC
Correlation with $\pi_{av,k}$	-0.390	Asymmetric TC
Correlation with distance coeff	-0.287	Symmetric TC
Correlation with $\pi_{av,k}$	-0.400	Symmetric TC

Notes: Excluding services; coefficients are all significant at 5%.

We also examine the correlation between income elasticity and an index inversely related to “tradability” $\pi_{av,k} = \frac{1}{N} \sum_n \pi_{nnk}$ which corresponds to the average share of consumption of domestic goods for sector k ($\pi_{av,k} = \frac{1}{N}$ if trade is costless). This term plays a role in the decomposition for trade-to-GDP ratios described in sections 2.4.2 and 4.2. As shown in Table A2, we find it to be significantly correlated with income elasticity across goods, which implies that high-income countries tend to consume goods that are more likely to be traded.

Skill- and Theta-driven Models

These models are estimated in one step by regressing final demand trade flows as described in Equations (27) and (29). The pseudo-R2 (squared correlation between observed and fitted trade) of the Theta-driven Model is 0.972 (0.404 if excluding domestic demand) whereas the pseudo-R2 of the Skill-driven Model is 0.985 (0.612 if excluding domestic demand). We obtain a better fit with the Skill-driven Model than the Theta-driven Model. In comparison, the R2 of the unconstrained model using both importer and exporter fixed effects to regress final demand trade flows yields an R2 of 0.999 (0.863 if excluding domestic demand).

B) A reduced-form approximation for the second step

In this appendix section, we propose a reduced-form approach to estimate income elasticities which approximates our second-step structural estimation of “Constant Relative Income Elasticities” preferences. There is no closed-form solution for λ_n as a function of per capita income e_n except in the homothetic case where $\sigma_k = \sigma$ is constant across goods. Nevertheless, we show in Table 1 that there is a -98% correlation between $\log \lambda_n$ and $\log e_n$. If we approximate $\log \lambda_n \approx \rho_1 \log e_n + \rho_0$ as a linear function of \log per capita income, the second-stage

specification can be rewritten as:

$$\log x_{nk} \approx \alpha_{7,k} + \rho_1 \sigma_k \log e_n + \mu_k \Phi_{nk} + \varepsilon_{nk}$$

where $\alpha_{7,k} = \rho_0 + \log \alpha_{5,k}$ is an industry fixed effect. This specification resembles AIDS (Deaton and Muellbauer, 1980) with the difference that the dependent variable is the log of expenditures instead of consumption shares. In this specification, we can retrieve the sigma's up to a constant: $\widehat{\rho_1 \sigma_k}$. We can then use fitted expenditures and these estimates to retrieve an approximation of income elasticities:

$$\eta_{nk} \approx \widehat{\rho_1 \sigma_k} \frac{\sum_{k'} \hat{x}_{nk}}{\sum_{k'} \hat{x}_{nk'} \widehat{\rho_1 \sigma_{k'}}$$

This approximation of income elasticities should not depend on the ρ_1 coefficient. We implement this specification and find a 98 % Spearman-rank correlation with our benchmark structural estimates, showing that this reduced-form approach provides a very good approximation (see Figure A1).

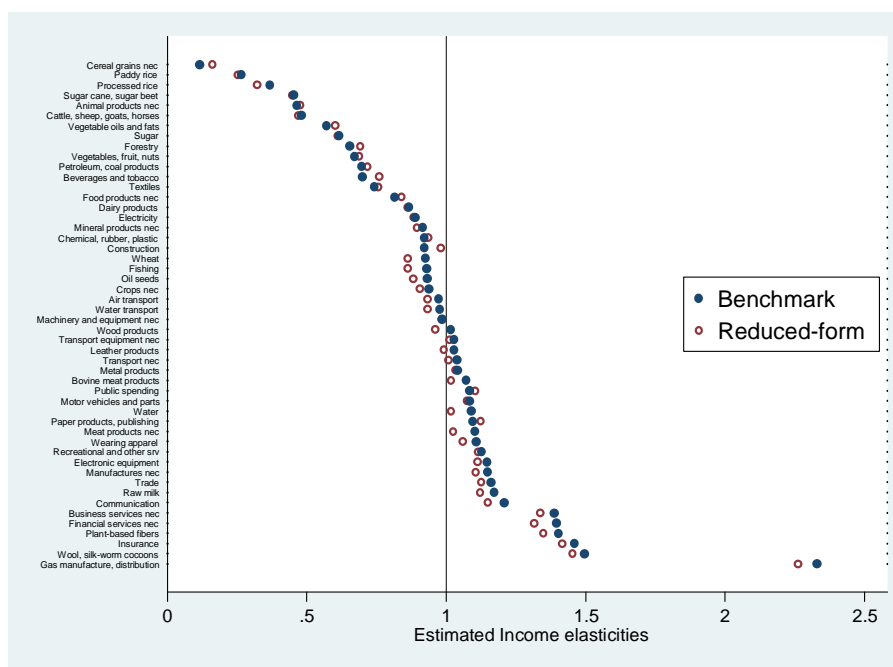


Figure A1: Benchmark vs. reduced-form approximation

C) Using actual price data

In Section 3, income elasticities are estimated by controlling for supply-side characteristics using a proxy price index P_{nk} which is constructed from the estimated Φ_{nk} (from the gravity

equations). Possible mis-estimation of this unobserved variable might raise concerns that our income elasticity estimates are biased. To test for this, we use actual price data from the 2005 International Comparison Program (ICP) (World Bank 2005), an extensive dataset which includes price indices for a wide range of products and countries. Despite mapping issues, we are able to match ICP price indices to 38 of the 50 sectors and 88 out of 94 countries included in our analysis.

The idea here is not to test whether the estimated P_{nk} perfectly match the actual price indices, as there are many reasons for them not to. Indeed, a regression of the log of the ICP price index on $\log P_{nk}$ including both country and sector fixed effects reveals a significant but weak correlation (beta coefficient = 0.136, p-value < 0.001).

Rather, we are interested in knowing if the inclusion of ICP price data in demand estimations leads to significantly different income elasticity estimates. A simple reduced-form log regression of final demand D_{nk} on both price indices (not shown), with both region and sector fixed effects, reveals that the constructed P_{nk} have a stronger explanatory power than the ICP index (beta coefficient of -0.213 versus -0.044, both p-values < 0.01).

Including the ICP price index in the estimation of CRIE demand parameters in our benchmark specification confirms that its predictive power is less than that of the constructed P_{nk} . Indeed, resulting income elasticity estimates are closer to those obtained by ignoring prices entirely. Table A3 displays the correlation of income elasticity and skill intensity when income elasticities are estimated using ICP prices (column 2) and using both indices (column 3). We clearly find that controlling for supply-side characteristics with our proxy price index P_{nk} has a greater impact on demand estimates. Thus, without being a definite test of the validity of our price index proxy, the comparison with external price data suggests that potential mis-estimation of the Φ_{nk} would tend to bias our correlation estimates downwards, if anything.

D) Alternative demand systems

In order to test how our CRIE income elasticity estimates stack up against other demand systems, we compare them with estimates - generated using the same dataset - from two well-known alternative demand systems which also exhibit non-homothetic behavior: the linear expenditure system (LES) and the "Almost Ideal Demand System" (AIDS). LES is derived from Stone-Geary preferences and is essentially an origin-displaced Cobb-Douglas function. AIDS, first introduced by Deaton and Muellbauer (1980), is not derived from any particular utility function, but has been widely used for its aggregation properties and its simplicity. Under the assumption of identical relative prices across regions, these demand systems can be shown to yield the following relationship between sectoral consumption shares and per-capita

Table A3: Skilled labor to income elasticity correlation - Robustness across specifications

Demand system:	CRIE			LES	AIDS	
Dependent variable:	Log expenditure		Expenditure shares	Expenditure shares	Expenditure shares	
Prices:	Phi (bmk)	ICP	Both	-	-	-
Region(s):	(1)	(2)	(3)	(4)	(5)	(6)
All	0.523	0.704	0.682	0.819	0.668	0.858
With robust data	0.488	0.638	0.632	0.766	0.614	0.805
USA	0.396	0.528	0.502	0.629	0.394	0.630
EU	0.524	0.579	0.603	0.703	0.580	0.719
Japan	0.466	0.590	0.581	0.760	0.748	0.829
Observations	50	38	38	50	50	50

Notes: all income elasticities calculated using median country expenditure shares. All correlations are significant at the 1% level.

expenditures:

$$\text{LES} : \frac{x_{nk}}{\sum_k x_{nk}} = \alpha_k + \gamma_k e_n^{-1} \quad \text{AIDS} : \frac{x_{nk}}{\sum_k x_{nk}} = \alpha_k + \gamma_k \log e_n$$

Note that the budget constraint imposes $\sum_k \alpha_k = 1$ and $\sum_k \gamma_k = 0$ in both cases. In each case, this relationship is estimated by sector by minimizing squared errors in expenditure shares. For the sake of the comparison, we also reestimate CRIE preferences by minimizing errors in expenditure shares (whereas our benchmark estimates minimize errors in log expenditures). The resulting estimates of α_k and γ_k are then used to compute income elasticities η_{nk} with LES and AIDS as:

$$\text{LES} : \eta_{nk} = \alpha_k (\gamma_k + \alpha_k e_n^{-1})^{-1} \quad \text{AIDS} : \eta_{nk} = 1 + \gamma_k (\alpha_k + \gamma_k \log e_n)^{-1}$$

Figure A2 plots the distribution of these income elasticities against the CRIE estimates. All estimates are evaluated using median-income country per-capita expenditures. Clearly, CRIE estimates are in line with both of these alternative demand systems. Spearman coefficients of rank correlation with CRIE estimates are 0.88 for LES and 0.85 for AIDS. Most importantly, columns (5) and (6) of Table A3 confirm that the result of strong correlation between income elasticities and skill intensity is robust across all three demand systems. Comparing the fit of each demand system confirms that CRIE preferences (R2= 0.782) and AIDS (R2=0.774) describe consumption patterns better than LES (R2=0.752). In comparison, fixed effects alone yield a R2 of 0.719.

Figure A2 also reveals the weakness of the LES demand system: income elasticities are very sensitive to income and converge rapidly to unity as income increases. Thus, even when evaluated at the median country income (as in Figure A2), income elasticities exhibit small deviations to one. AIDS performs better and yields a larger variability which is closer to that generated by CRIE. However, AIDS consumption patterns cannot be derived from actual preferences, as noted by Deaton and Muellbauer (1980).

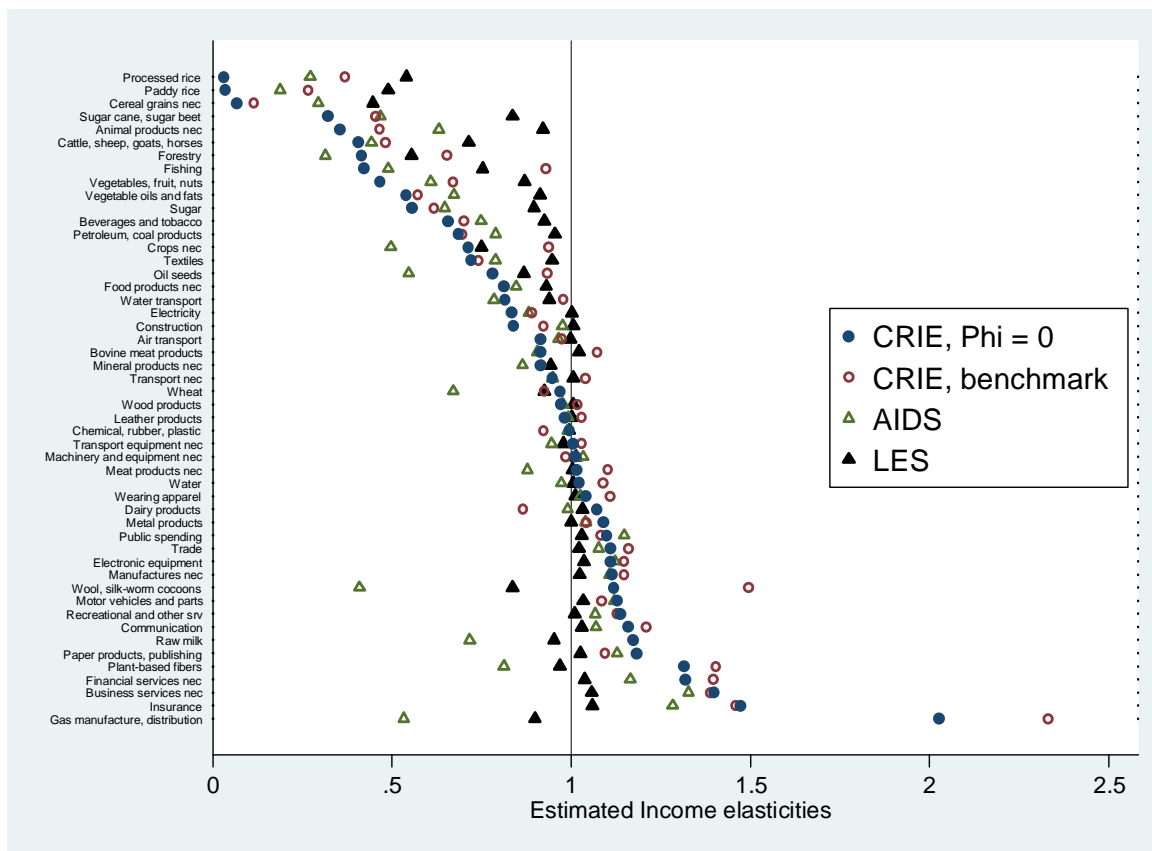


Figure A2: Comparison of distribution of income elasticities across demand systems

E) Measurement of skill intensity

All results from the previous sections are estimated using sectoral skill intensity indices which are computed as an average over all countries. We now test whether the main correlations of interest are robust to using skill intensity measured on the subset of countries with the most reliable data (see Section 3.3). Table A3 displays the correlation of income elasticities to skill intensity using different regional subsets of the GTAP data : all GTAP regions, the reliable

regions, the US, Europe (EU) and Japan. Although, the correlation seems to generally be smaller for the USA than for the EU and Japan, it remains large and significant for all regions.

F) Within-country income distribution

Compared to a hypothetical situation where income is homogeneous within each country, within-country income inequalities can reduce the observed variations in consumption patterns between countries. For instance, income inequalities could explain why it is possible to find luxurious cars in Africa (as other high-income elastic goods), while this type of goods is clearly not purchased by any individual with the average income of an African country.

Empirically, income inequalities create a downward bias in the dispersion of estimated income elasticities if they are not taken into account (i.e. it biases our income elasticity estimates towards unity). Conversely, accounting for within-country inequalities should *reinforce* the differences in estimated income elasticities across sectors: it would be otherwise difficult to explain the large differences in observed consumption patterns across countries.

To confirm this intuition and test the sensitivity of our results to the inclusion of within-country inequalities, we rely as in Fielser (2011) on World Bank data describing the share of total income held by different percentiles of the population. Available data covers 7 income classes (the first two and last two deciles, as well as the 3 middle quintiles) and 89 of the 94 countries in our sample.

The estimation procedure in Section 3 is modified to allow for seven representative consumers in each country (each representing one population quintile or decile). As expected, resulting income elasticity estimates exhibit slightly larger variations across sectors, although the differences are very small: the mean absolute deviations from one increases from 0.261 to 0.270.⁴⁵ The correlation between the two series is 0.995.

Correspondingly, the correlation of these income elasticity estimates with skill intensity increases only very slightly (from 0.523 to 0.528). Thus, accounting for within-country dispersion in incomes increases the estimated effect of non-homotheticity and makes our results stronger. However, given the small magnitude of the bias, we are comfortable with using the estimates from Section 3 for counterfactual analyses.

G) Factor content of trade: accounting for traded intermediate goods

To better account for trade in intermediate goods and differences in factor requirement matrices across countries, we compute the factor content of trade following the method developed by

⁴⁵comparing estimates generated with the comparable set of 89 countries

Trefler and Zhu (2010). This computation relates to rows 5 and 7 of Table 4 in Section 4.1 on the “missing trade” puzzle.

As in Trefler and Zhu (2010), we construct the matrix A_{kfn}^P of direct and indirect factor requirements by taking into account factors embodied in traded intermediate goods. Data on domestic and imported input requirements at the country level are provided in the GTAP database. To further approximate bilateral vertical linkages between any two countries, the construction of the matrix A^P relies on a proportionality assumption (see below). The factor content of trade is then constructed as $F_{nf} = A_{kfn}^P Y_{nk} - \sum_i A_{kfi}^P X_{nik}$. We further compute A_{kfn}^D as the matrix of factors embodied in consumption of final goods consumed by country n . Assuming that the share of final goods purchased from source i equals the share of imports from this source (proportionality assumption), we define A^D as:

$$A_{kfn}^D = \frac{\sum_i A_{kfi}^P X_{nik}}{\sum_i X_{nik}} \quad (32)$$

Building on Lemma 1 of Trefler and Zhu (2010), we show (see proof below) that the factor content of trade satisfies:

$$F_{nf} = \left[w_{nf} V_{nf} - s_n \sum_i w_{if} V_{if} \right] - \sum_k \left[A_{kfn}^D D_{nk} - s_n \left(\sum_i A_{kfi}^D D_{ik} \right) \right] \quad (33)$$

This is an accounting equality, which is exactly satisfied by construction when we use observed demand D_{nk} and A_{kfn}^D to construct the right-hand-side term. In the results presented in Table 4, we construct the predicted factor content of trade using fitted demand \widehat{D}_{nk} and fitted factor content of consumption \widehat{A}_{kfn}^D to illustrate the role of non-homotheticity and trade costs. In particular, when we assume homothetic preferences and no trade costs, the right-hand side reduces to $w_{nf} V_{nf} - s_n \sum_i w_{if} V_{if}$, even if the factor content coefficients A_{kfn}^P vary across countries. This corresponds to the “consumption similarity” condition emphasized by Trefler and Zhu (2010). The second term reflects the correction needed when the consumption similarity condition is violated, which is the case since trade is not costless and preferences are not homothetic.

Measurement of FTC with intermediate goods: matrix \mathbf{A}^P

In their definition of the factor content of trade, Trefler and Zhu (2010) construct a matrix \mathbf{A}^P reflecting the factor content of final goods production by taking into account the factor content of traded intermediate goods.

In line with Trefler and Zhu (2010), we define the trade matrix \mathbf{T} as:

$$\mathbf{T} = \begin{bmatrix} (\sum_{n \neq 1} \mathbf{X}_{n1}) & -\mathbf{X}_{21} & \cdots & -\mathbf{X}_{N1} \\ -\mathbf{X}_{12} & (\sum_{n \neq 2} \mathbf{X}_{n2}) & \cdots & -\mathbf{X}_{N2} \\ \vdots & \vdots & \ddots & \vdots \\ -\mathbf{X}_{1N} & -\mathbf{X}_{2N} & \cdots & (\sum_{n \neq N} \mathbf{X}_{nN}) \end{bmatrix}$$

where each \mathbf{X}_{ni} is the vector of trade flows X_{nik} for goods k shipped from i to n :

$$\mathbf{X}_{ni} = \begin{bmatrix} X_{ni1} \\ X_{ni2} \\ \vdots \\ X_{niK} \end{bmatrix}$$

we define the matrix \mathbf{D} for trade in final goods as:

$$\mathbf{D} = \begin{bmatrix} \mathbf{D}_{11} & \mathbf{D}_{21} & \cdots & \mathbf{D}_{N1} \\ \mathbf{D}_{12} & \mathbf{D}_{22} & \cdots & \mathbf{D}_{N2} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{D}_{1N} & \mathbf{D}_{2N} & \cdots & \mathbf{D}_{NN} \end{bmatrix}$$

where each \mathbf{D}_{ni} is the vector of trade flows D_{nik} for final goods k shipped from i to n . Given our proportionality assumption, we measure D_{nik} as $\frac{D_{nk}X_{nik}}{\sum_i X_{nik}}$. We also define the production matrix \mathbf{Y} as the diagonal matrix:

$$\mathbf{Y} = \begin{bmatrix} \mathbf{Y}_1 & 0 & \cdots & 0 \\ 0 & \mathbf{Y}_2 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \mathbf{Y}_N \end{bmatrix}$$

where each \mathbf{Y}_n is the vector of production Y_{nk} of country n in sector k .

The input-output matrix is denoted by \mathbf{B} and reflects the use of inputs k from a particular source i for each downstream industry k' and destination country n . In a matrix form:

$$\mathbf{B} = \begin{bmatrix} \mathbf{B}_{11} & \mathbf{B}_{21} & \cdots & \mathbf{B}_{N1} \\ \mathbf{B}_{12} & \mathbf{B}_{22} & \cdots & \mathbf{B}_{N2} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{B}_{1N} & \mathbf{B}_{2N} & \cdots & \mathbf{B}_{NN} \end{bmatrix}$$

where each individual block \mathbf{B}_{ni} is a K by K matrix where each cell $B_{nikk'}$ reflects the use of intermediate goods k' from country i for the production of one dollar of good k in country n .

We do not have data on each component $B_{nikk'}$, as GTAP only provides information on domestic input uses $B_{nnkk'}$ and aggregate imported intermediates $B_{nkk'}^{imp}$ for each downstream industry k' , upstream industry k in country n . Assuming that import shares from a particular source i do not depend on the downstream industry k' (proportionality assumption), we construct $B_{nikk'}$ as:

$$B_{nikk'} = B_{nkk'}^{imp} \frac{X_{nik}}{\sum_{i \neq n} X_{nik}}$$

As goods are either used as final goods or intermediate goods, we obtain the following accounting equality described in Trefler and Zhu (2010):

$$\mathbf{T} = \mathbf{Y} - \mathbf{B}\mathbf{Y} - \mathbf{D}$$

which leads to:

$$\mathbf{Y} = (\mathbf{I} - \mathbf{B})^{-1}(\mathbf{T} + \mathbf{D}) \quad (34)$$

We now define matrix \mathbf{A}^P as:

$$\mathbf{A}^P = \Xi(\mathbf{I} - \mathbf{B})^{-1}$$

where Ξ is the matrix of direct factor requirements, i.e. the matrix of coefficients β_{kfn} reflecting the value of factor f required in the production of goods k in country n . Since the value of factor k used in production in country n is given by $w_{nf}V_{nf} = \sum_k \beta_{kfn}Y_{nk}$, we obtain from Equation (34) that $\mathbf{A}^P(\mathbf{T} + \mathbf{D})$ equals the matrix of payments to factors (i.e. each $\mathbf{A}^P(\mathbf{T}_n + \mathbf{D}_n)$ is the vector of payments to factors used in country n).

Trefler and Zhu (2010) then defines the factor content of trade by taking the trade component of the above equation: $\mathbf{F} = \Xi(\mathbf{I} - \mathbf{B})^{-1}\mathbf{T}$ (where each component $\mathbf{A}^P\mathbf{T}_n$ is the vector of factor content of trade for country n for each column-vector of trade \mathbf{T}_n).

Factor content of trade: proof of Equation (33)

To prove Equation (33), we first show that:

$$\sum_k A_{kfn}^D D_{nk} + F_{nf} = w_{nf}V_{nf}$$

and then show that:

$$\sum_{k,n} A_{kfn}^D D_{nk} = \sum_n w_{nf}V_{nf}$$

We obtain Equation (33) by combining these two equations.

To see the first equality, we can write:

$$\begin{aligned} F_{nf} + \sum_k A_{kfn}^D D_{nk} &= F_{nf} + \sum_{k,i} A_{kfi}^P D_{nik} \\ &= \sum_{k,i} A_{kfi}^P (T_{nik} + D_{nik}) \end{aligned}$$

where the first line is obtained from the definition of A^D and the second line from the definition of the factor content of trade F_{nf} . In matrix form, we know however that $\mathbf{Y} = (\mathbf{I} - \mathbf{B})^{-1}(\mathbf{T} + \mathbf{D})$ (equation 34) and that the matrix \mathbf{A}^P is defined as $\mathbf{A}^P = \Xi(\mathbf{I} - \mathbf{B})^{-1}$ where Ξ is the matrix with factor requirement coefficient β_{kfn} . Hence: $\mathbf{A}^P(\mathbf{T} + \mathbf{D}) = \Xi\mathbf{Y}$ and we thus obtain:

$$\sum_{k,i} A_{kfi}^P (T_{nik} + D_{nik}) = \sum_k \beta_{kfn} Y_{nk} = w_{fn} V_{fn}$$

We now prove the second equation mentioned earlier, i.e. that $\sum_{k,n} A_{kfn}^D D_{nk} = \sum_n w_{nf} V_{nf}$. Using the definition of A^D , we can write:

$$\sum_{k,n} A_{kfn}^D D_{nk} = \sum_{k,n,i} A_{kfi}^P D_{nik}$$

Given that the sum of the coefficients of matrix \mathbf{T} equals zero for each row, we can also write:

$$\sum_{k,n} A_{kfn}^D D_{nk} = \sum_{k,n,i} A_{kfi}^P (D_{nik} + T_{nik})$$

(where the T_{nik} denote the coefficients of matrix \mathbf{T}). As described above, $\sum_{k,i} A_{kfi}^P (D_{nik} + T_{nik})$ equals the payment to factor f in country n . Hence this proves the second equality:

$$\sum_{k,n} A_{kfn}^D D_{nk} = \sum_n w_{nf} V_{nf}$$