Wage Inequality and Skill Supplies in a Globalised World

Lorenzo Rotunno
Adrian Wood

Wage inequality and skill supplies in a globalised world[*]

Lorenzo Rotunno[†]

Adrian Wood[‡]

This draft: November 2017

Abstract

We investigate empirically, and explain theoretically, how the relative wages of skilled and unskilled workers vary with their relative supplies in open economies. Our results resolve a conflict between the predictions of standard trade theory and experience of how labour markets work. We show that relative wages respond to variation in relative skill supplies in countries that trade, as intuition and much other evidence suggest, but also that the wage response decreases as trade barriers fall and that, as trade theory suggests, is weak in very open economies.

Keywords: Heckscher-Ohlin, trade and wages, wage inequality, labour markets.
JEL codes: F11, F16, J23, J31.

[†]Aix-Marseille Univ., CNRS, EHESS, Centrale Marseille, AMSE. Chemin du Château Lafarge, 13290 Les Milles, France. Office B987/988; Tel.: +33 (0)442935987. Email: lorenzo.rotunno@univ-amu.fr

[‡]Department of International Development, University of Oxford. Email: adrian.wood@qeh.ox.ac.uk

I Introduction

A commonly suggested policy response to widespread concerns about wage inequality between skilled and unskilled workers is more education and training to increase the relative supply of skilled workers. However, and even abstracting from the practical and financial difficulties involved, standard economic models do not offer a unified view on whether an increase in supply would reduce this inequality in economies open to international trade. In a one-cone Heckscher-Ohlin (HO) trade model, wages in an open economy are pinned down by world prices and trade costs, and would not be altered by a change in skill supplies. Labour economists, though, tend to doubt that "your wages are set in Beijing" (Freeman, 1995), and would expect an increase in the relative supply of skilled workers in an open economy to reduce their relative wage, as in a closed economy.

In more sophisticated trade models, this sharp divergence of predictions is blurred. Multi-cone HO theory (Markusen and Venables, 2007; Leamer, 2012) allows countries with widely differing relative skill supplies to have different wages because they produce different goods. Firm heterogeneity and labour market frictions (Helpman and Itskhoki, 2010) can also link relative wages with relative skill supplies, though not necessarily in the direction that labour economists would expect (Trionfetti, 2015). But there remains an underlying inconsistency between standard trade theory (which sees changes in production structure as the main absorber of shocks in skill supplies) and labour market intuition (which sees a natural relationship between relative wages and skill supplies). The tension emerges, for example, in Morrow and Trefler (2017, n. 13) and in the work of Burstein et al. (2017) on the impact of immigration on wages in tradable and non-tradable occupations.

Empirical resolution of this theoretical disagreement has been hindered by a shortage of panel data on wages and skill supplies in a wide spectrum of countries. A big step forward in this respect recently is the Socio-Economic Accounts of the World Input-Output Database (WIOD), which cover 40 countries during 1995-2009. What these data show at first sight is set out in Figure 1, which plots the wages of college-educated relative to other workers against their relative employment, both in logs, after averaging the data over time (Figure 1(a)) and in changes over the period (Figure 1(b)) – more information on the variables is in section III. There is a clear downward slope, with an elasticity precisely estimated at -0.27 in the averaged data (and consistently between -0.2 and -0.3 in cross-section regressions by year). The relationship in changes is confused by outliers, but if we drop Mexico, whose employment data are suspect,[1] and the transition countries, whose labour market institutions

---

[1]The fall in the college-educated share of employment in Mexico is out of line with all other countries,

changed substantially, there is a similar, though less significant, elasticity of -0.26.

Figure 1: Relative wages and skill supplies across countries

Figure 1(a): Averaged data, 1995-2009

Figure 1(b): Changes 1995-2009
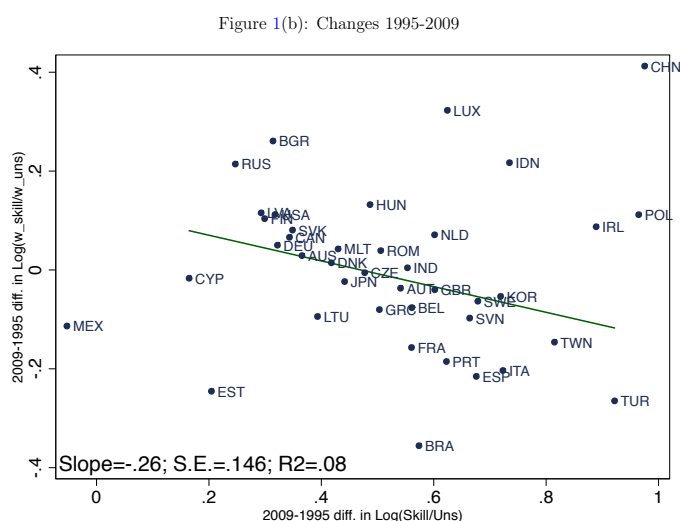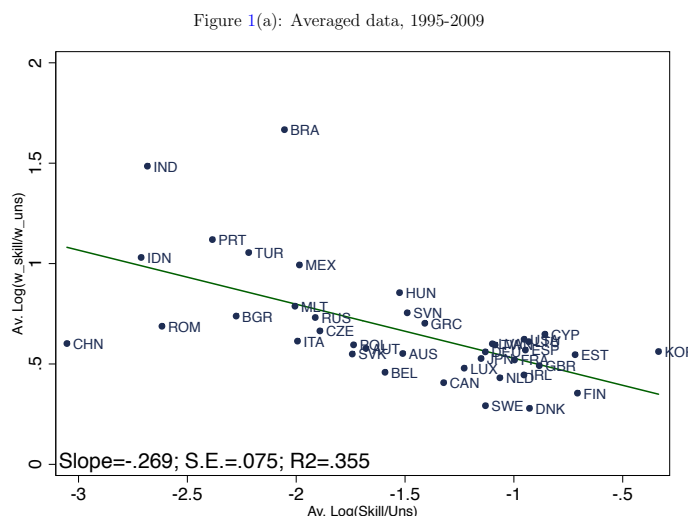
Notes: All regressions include a constant. Prediction line and statistics in Figure 1(b) are based on a regression excluding CHN, MEX, and POL.

These numerical values invite thought. Their downward slope is inconsistent with the zero predicted by a one-cone HO model, but far less steep than the -0.7 implied by the often-quoted Katz and Murphy (1992) aggregate elasticity of substitution of 1.4 between college-educated and other labour in the US.[2] The cross-country slope can be converted, as

in section A.III of the Online Appendix, into a step function consistent (apart from China) with a four-cone HO model.[3] However, half of the countries stay in one cone throughout the period, so in theory their wages should have been insensitive to changes in skill supplies, but there is no obvious clustering of relative wage changes around zero in Figure 1(b).

In this paper, we seek to explain why the observed elasticity is of this sign and magnitude, and in the process to narrow the gap between standard trade theory and experience of how labour markets work. Empirically, we show that relative wages in open economies vary with relative skill supplies, as labour economists believe, but that the response of wages to skill supplies is smaller in countries with lower barriers to trade, and in very open economies comes close to the simplest trade economist view that wages are unaffected by endowments. Theoretically, we explain this pattern by extending the Heckscher-Ohlin-Samuelson (HOS) model to include elements of other more recent models.

A distinguishing characteristic of our extended Heckscher-Ohlin (EHO) model, compared to HOS, is that from the perspective of an individual country, demand and supply in world markets are less than infinitely elastic. Absorbing changes in skill supplies by changes in output mix thus requires changes in goods prices. A longstanding explanation of inelasticity, due to Armington (1969) and embodied in many Computable General Equilibrium (CGE) and gravity models (Anderson, 1979), as well as in Krugman (1979), is qualitative differences among the varieties of goods made by different countries or firms. More recently, however, inelasticity has emerged as a feature of models with stochastic variations in efficiency among countries (Eaton and Kortum, 2002) and firms (Melitz, 2003). As shown by Arkolakis et al. (2012) and Costinot and Rodriguez-Clare (2014), these two sources of inelasticity are on certain assumptions interchangeable.

Imperfect substitutability of home and foreign varieties for Armington or Eaton-Kortum reasons carries over into our model by creating inelasticity of demand in labour markets and thus helping to explain why relative wages vary with skill supplies in open economies. Our model, however, includes two other sources of labour demand inelasticity (drawn from Wood (2012)). First, the goods price elasticities faced by firms depend also on the degree of

and Autor (1999) in their review of the empirical evidence from different countries in the 1990s. For later years in the US, Blankenau and Cassou (2011) also estimate an elasticity of about -0.7, as do Acemoglu and Autor (2011, Table 6), though only with a non-linear time trend. Blum (2010) too estimates an elasticity of about -0.7 for 15-year changes in all countries, using data for 1973-90 on occupational (production vs non-production workers) rather than educational categories (with less negative values for shorter periods and more negative values for developed countries only).

[3]Most empirical estimates of multi-cone HOS models analyse product mix or choice of technique (Davis and Weinstein, 2001; Schott, 2003). Kiyota (2011, 2012) also analyses factor price variation, but only within one country, as in most studies of factor price equalisation (Hanson and Slaughter, 2002; Tomiura, 2005; Bernard et al., 2008, 2013).

openness of the economy: as was noted by Johnson and Stafford (1999), if home suppliers have larger shares of the home market, as a result of high trade barriers, their sales become more affected by elasticities of substitution among goods (which are lower than those among varieties of goods). Second, the existence of other costs that do not vary in proportion to labour costs (per-unit trade costs, for example), and thus increase the size of the relative wage changes needed to achieve price-induced skill-supply-absorbing changes in the composition of output.

The relative elasticity of demand for skilled and unskilled workers in an open economy therefore depends on all three of these forces – substitutability between home and foreign varieties, trade shares, and the size of costs that do not vary in proportion to labour costs. This paper sets out an EHO model that includes these three forces and provides the basis for our empirical investigation of the responsiveness of relative wages to variation in skill supplies using the WIOD data.

Reduced-form estimation of the relationships involved is problematic: observed elasticities of the sort in Figure 1 are the product of several forces, including the three just mentioned, and there is not enough exogenous variation in the data to be able to identify their influence separately. With this caveat in mind, we nonetheless start by using reduced-form methods to check on the prima facie consistency of our EHO model with reality. As expected from Figure 1, our reduced-form results confirm much other evidence (e.g. Katz and Murphy (1992); Robbins (1996); Blum (2010); Marshall (2012); Morrow and Trefler (2017)) of an inverse relationship between the relative wages and relative supplies of skilled and unskilled workers. As predicted by our model, moreover, the inverse relationship is weaker where countries are more open to trade and where the ratio of wages to non-proportional trade and production costs is higher. For countries in the upper quartile of trade openness, the inverse relationship disappears.

These inferences from the reduced-form results are supported by structural estimation of the EHO relationship between the relative wages and relative supplies of skilled workers. Predicted changes in relative wages match the variation in the data, both across countries and within countries over time, and the fit of the predictions is improved by incorporating the two new forces in the model - trade shares and non-proportional cost wedges.

The structural estimates reveal, however, that the demand-elasticity-reducing features of EHO predict a wage-skill supply elasticity that is much lower (i.e. more negative) than the reality of Figure 1 and our reduced-form estimates. The difference can plausibly be explained by two other elements that we add to our baseline EHO model, drawing on Acemoglu (2007) and Caron et al. (2014), who identify mechanisms by which an increased relative

supply of skill shifts the relative demand for skill upwards. One is directed technical change: increases in the supply of a factor create an incentive to find new ways of using it more productively. The other is that increases in skill supply increase income and that the demand for skill-intensive goods is income-elastic.

Our paper identifies theoretically and quantifies empirically the role of output and sectoral skill intensity responses to variation in factor supplies in determining the wage-skill supply elasticity. Romalis (2004) and Chor (2010) provide evidence for the basic HO process of absorption of variation in factor supplies by variation in the sectoral composition of trade. Relatedly, Lafortune et al. (2015) find that Rybczynski effects on output composition help to absorb immigration-driven changes in labour supply. Unlike HOS, however, the changes in output composition in our EHO model are induced by changes in relative goods prices, as in the multi-factor multi-sector analysis of Costinot and Rodriguez-Clare (2014, 221-3), which we extend by making goods price elasticities vary with a country's degree of openness (in the spirit of Rodrik (1997)).[4] Non-homothetic demand reduces the required goods price changes, as in Caron et al. (2014).

Variation in sectoral output structure is far from sufficient to absorb all the variation in skill supplies, as documented by Blum (2010), Ciccone and Peri (2005) and much of the migration-related literature (e.g. Lewis, 2013). The rest is absorbed by variation in sectoral skill intensity, as in Davis and Weinstein (2001) and Morrow and Trefler (2017), induced in EHO by the relative wage changes that are caused by changes in relative goods prices, and reinforced by directed technical change, as in Acemoglu (2007). The effect of relative goods price changes on relative wages is amplified by non-proportional trade and other costs, as in Alchian and Allen (1964) and Hummels and Skiba (2004). Observed changes in sectoral skill intensity reflect changes in intra-sectoral output structure, as well as in production techniques (Davis and Weinstein, 2001; Schott, 2003).

Because relative wages vary with relative skill supplies in EHO, there cannot be factor price equalisation, even in relative productivity-adjusted terms. There are many other reasons for cross-country differences in relative wages, including the Ricardian differences in technology emphasised by Morrow and Trefler (2017), and differences in market size as in Epifani and Gancia (2006), but most of them are not addressed in this paper because

---

[4]Rodrik showed theoretically that greater international mobility of capital could increase the elasticity of demand for labour in an economy open to trade, and he conjectured that an increase in openness to trade would have a similar effect (Panagariya, 1999). His conjecture has been tested at plant, firm and industry level (but not for the whole economy) in later studies, most of which have found some relationship between these micro elasticities, especially for unskilled labour, and proxies for openness to trade, though in some cases not robust to controls for new technology or time trends (e.g., Slaughter (2001); Hijzen and Swaim (2010); Senses (2010)).

they do not affect the wage-skill supply elasticity that is our focus. An exception is the influence of trade on relative wages, but even in this regard we focus on the effects of openness on the elasticity of the wage-skill supply relationship and neglect most other ways in which international linkages can make a country's relative wages higher or lower than would be expected from its relative skill supply (e.g. Wood (2002); Epifani and Gancia (2008); Grossman and Rossi-Hansberg (2008); Harrison et al. (2011); Haskel et al. (2012); Burstein and Vogel (2016)).

The rest of the paper is organised as follows. Section II sets out the theory. Section III describes our empirical strategy and the WIOD data. Sections IV-VI present the results. Section VII concludes.

II Extended Heckscher-Ohlin

To explain our argument in a simple way, drawing on Wood (2012), we set out a 2×2 neoclassical model of a single country in the small-proportional-changes 'hat' algebra of Jones (1965). Two factors, H (high-skilled workers) and L (low-skilled workers), produce two final goods, B (biochemicals, which are H-intensive) and G (garments, which are L-intensive). Later, we include an unlimited number of goods and discuss the roles of intermediate inputs, capital and other influences on the relative wage-skill supply elasticity.

With given technology, changes in the relative labour costs of the goods, c, are related to changes in relative wages, w, by

$$(1) \quad \widehat{c}_{B/G} = (\theta_{HB} - \theta_{HG}) \, \widehat{w}_{H/L}$$

where θ_{ij} is the share of labour type i in the cost of good j, $\widehat{x} = {}^{dx}/_{x}$, and $\widehat{x}_{1/2} = \widehat{x}_1 - \widehat{x}_2$. Since $\theta_{HB} > \theta_{HG}$, a rise in the relative wage of skilled workers causes a rise in the relative cost and price of the skill-intensive good. Labour-market clearing requires

$$(2) \quad \widehat{v}_{H/L} = - \left[1 - (\lambda_{HB} - \lambda_{LB}) \, (\theta_{HB} - \theta_{LB}) \right] \sigma \widehat{w}_{H/L} + (\lambda_{HB} - \lambda_{LB}) \widehat{q}_{B/G}$$

where the economy-wide supply of a labour type is denoted by v, the output of a good by q, λ_{ij} is the share of the supply of labour type i used by good j (so $\lambda_{HB} > \lambda_{LB}$), and σ is the elasticity of substitution between H and L in production, assumed to be the same for both goods. A rise (say) in the relative supply of skilled labour must be matched by a rise in the relative demand for skilled labour, which can be achieved by a fall in the relative skilled wage that induces a rise in the skill-intensity of the techniques used in producing

both goods (the first right-hand side (rhs) term in (2)) and/or by a shift in the composition of output towards the skill-intensive good B (the second term).

(a) Demand system

The relative sales of goods B and G by home producers depend on their relative prices because for each good there is a finite 'trade elasticity' that links the share of imports in domestic expenditure to the relative prices of its imported and home-produced varieties.[5] As in Arkolakis et al. (2012) and Costinot and Rodriguez-Clare (2014), we assume within each sector a CES utility function

$$(3) \quad C_j = \left[\left(C_j^H \right)^{\frac{\beta_j - 1}{\beta_j}} + \left(C_j^M \right)^{\frac{\beta_j - 1}{\beta_j}} \right]^{\frac{\beta}{\beta_j - 1}}$$

where C_j^H and C_j^M are composites of home-produced and imported varieties in sector j, and β_j equals one plus the trade elasticity (since the latter refers to the value rather than the volume of sales). This elasticity may reflect adjustments at either the intensive margin (more or less consumption of qualitatively different varieties, as in Armington (1969) and Krugman (1979)) or the extensive margin (purchases of identical varieties from different countries or firms, as in Eaton and Kortum (2002) and, partly, in Melitz (2003)). It is convenient to assume that the elasticities of substitution within the home-produced and imported composites are equal in each sector j to the elasticity of substitution between them, β_j, though this is not always so (Costinot and Rodriguez-Clare (2014, p. 244-6); Feenstra et al. (2017)).

The relative sales of the two goods depend on the relative prices of the B and G utility aggregates of home-produced and imported varieties, according to a higher-level CES utility function

$$(4) \quad C = \left[\alpha C_B^{\frac{\gamma_{BG} - 1}{\gamma_{BG}}} + (1 - \alpha) C_G^{\frac{\gamma_{BG} - 1}{\gamma_{BG}}} \right]^{\frac{\gamma_{BG}}{\gamma_{BG} - 1}}$$

where α is a preference parameter and γ_{BG} is the elasticity of substitution between the goods, which is assumed to be lower than β_B or β_G, the elasticities of substitution among varieties.

The elasticity of the relative sales of domestic producers of B and G with respect to

---

[5]Strictly speaking, the trade elasticity refers to the ratio of imported to home-produced varieties rather than to the share of imports in expenditure, and is defined with respect to changes in trade costs rather than in prices more generally (Costinot and Rodriguez-Clare, 2014, p. 201).

their relative purchaser prices is thus a weighted average of γ_{BG}, β_B and β_G, in which the weights depend on the shares of domestic producers in the markets concerned. Intuitively, if they have only small shares, as in export markets open to all comers, the price elasticity is close to that among varieties, β. As market shares rise, as in home markets protected from foreign competition by trade barriers, the relative price of the B and G aggregates becomes increasingly dependent on the prices of home-produced varieties, so the relative sales of home producers become increasingly dependent on consumer preferences among goods (γ) rather than among varieties (β). If home producers supply the entire market, as with a non-traded good, the only relevant elasticity is γ.

Formally, with (3) and (4) being CES, the average elasticity in any particular market can be written precisely, following Sato (1967), as a weighted harmonic mean (see section V and Wood, 2012, section 2.1). For purposes of exposition, this average can temporarily be written more simply, and omitting the market superscript on all its terms, as a weighted arithmetic mean

$$(5) \quad \epsilon_{BG} = s_{BG}\gamma_{BG} + (1 - s_{BG})\beta_{BG}$$

where β_{BG} is an average of β_B and β_G and s_{BG} is the country's average share of the sales of these goods in the market concerned. For each good in each market, relative market share, s_j, is a decreasingly positive transformation of the ratio of its sales to those of all foreign suppliers, which is determined through the CES demand system as

$$(6) \quad \frac{p_j q_j}{p_j^* q_j^*} = R \left(\frac{p_j}{p_j^*}\right)^{1-\beta_j} = R \left[\frac{c_j (1 + \tau_j)}{c_j^* (1 + \tau_j^*)}\right]^{1-\beta_j}$$

with p being purchaser price, R a measure of the selling country's economic size relative to that of the rest of the world, and τ_j the ratio of other production and trade costs per unit of output of good j, t_j, to unit labour costs (ULCs), c_j. Market share depends on country size and on comparative advantage (reflected in c_j/c_j^*), and it increases with the height of trade barriers to foreign suppliers (relative to those that affect home-country suppliers – which would be lower in the home market, but similar or higher in an export market).

The effect of relative purchaser prices on a country's relative outputs of B and G depends on the average ϵ_{BG} across all its markets, at home and abroad, weighted by the shares of its total sales in each market, which for simplicity and with a slight abuse of notation we will also label ϵ_{BG}. This average elasticity increases with its openness to trade for two reasons. Lower barriers to imports push s_{BG} down and thus increase ϵ_{BG} in its home market. Greater

access to foreign markets increases the share of exports in its output and so the weight in the average ϵ_{BG} of the higher ϵ_{BG}'s in its foreign markets (where s_{BG}'s are small). The demand system in the 2×2 model can thus conveniently be summarised by

$$(7) \quad \widehat{q}_{B/G} = -\epsilon_{BG}\widehat{p}_{B/G}$$

in which a country's relative sales of the two goods depends on their average relative purchaser prices, p, and on the average relative demand elasticity, across all markets.

(b) Price-ratio elasticity

The elasticity of the relative wage with respect to relative skill supplies in our model depends not only on the purchaser price elasticity ϵ_{BG} discussed above, but also on the elasticity of the relative purchaser prices of goods B and G to their relative unit labour costs (ULCs), $c_{B/G}$, which we call the 'price-ratio elasticity', δ_{BG}. This elasticity is a novel feature of our model and is implicitly assumed in most other models to be unity, though incomplete pass-through plays a similar role in models with imperfect competition (Krugman (1979); Dornbusch (1987); and, with heterogeneous firms, Berman et al. (2012)).

In practice, δ_{BG} is normally less than unity (the relative purchaser prices of goods vary by proportionally less than their relative ULCs), because the purchaser prices of goods contain other-cost wedges (OCWs) that do not vary in proportion to their ULCs, for example per-unit trade costs (Alchian and Allen, 1964; Hummels and Skiba, 2004) – more OCW elements will be discussed later. Using the definition $\tau_j \equiv \frac{t_j}{c_j}$ as before, the price-ratio elasticity $\delta_{BG} = \frac{\widehat{p}_{B/G}}{\widehat{c}_{B/G}}$ is determined approximately by

$$(8) \quad \delta_{BG} = \frac{1 + \eta_{BG}\tau_{BG}}{1 + \tau_{BG}}$$

where τ_{BG} is the geometric mean of τ_B and τ_G, and η_{BG} is the elasticity of $\frac{t_B}{t_G}$ with respect to $\frac{c_B}{c_G}$.

To understand equation (8), consider the expression $\frac{1}{1+\tau_{BG}}$, which is what (8) would become if OCWs were strictly independent of ULCs ($\eta_{BG} = 0$), for example because all cost wedges were per-unit, and is simply the average share of labour costs in the purchaser prices of B and G. The smaller this share, as a result of a larger τ_{BG}, the smaller the effect on relative purchaser prices of a proportional change in relative ULCs (just as, for example, with c_j half of p_j, a 10% rise in c_j, with no change in t_j, would raise p_j by only

5%). However, insofar as relative OCWs vary with relative ULCs, for example because some trade costs are ad-valorem, η_{BG} will be positive, tending to increase δ_{BG} (and if η_{BG} were unity, as if for example OCWs consisted only of ad-valorem trade costs, δ_{BG} would be unity, too).

(c) Wage-skill supply elasticity in the 2×2 model

Combining equations (1), (2), and (7) (with $\widehat{p}_{B/G} = \delta_{BG}\widehat{c}_{B/G}$), the effect of changes in the relative supply of skilled and unskilled workers on relative wages can be derived as[6]

$$(9) \quad \widehat{w}_{H/L} = -\frac{1}{\left[1 - (\lambda_{HB} - \lambda_{LB})(\theta_{HB} - \theta_{HG})\right]\sigma + (\lambda_{HB} - \lambda_{LB})(\theta_{HB} - \theta_{HG})\epsilon_{BG}\delta_{BG}}\widehat{v}_{H/L}$$

The first term in the denominator of the rhs ratio is an 'aggregate' elasticity of substitution between H and L and shows how changes in wages induce supply-absorbing substitution between skill categories within sectors. The second term shows how changes in wages alter goods prices in ways that shift the sectoral composition of output in a direction that helps to absorb changes in skill supplies. It is the product of four elasticities: of relative labour costs with respect to relative wages $(\theta_{HB} - \theta_{HG})$, of relative goods prices with respect to relative labour costs δ_{BG}, of relative outputs with respect to relative goods prices ϵ_{BG}, and of relative employment of skilled and unskilled workers with respect to relative outputs $(\lambda_{HB} - \lambda_{LB})$.

The denominator of (9) is a weighted average of the elasticity of substitution parameter in production (σ) and the elasticity of relative sales with respect to labour costs $(\epsilon_{BG}\delta_{BG})$. The weights depend on the difference in skill intensity between the goods: the greater is $(\lambda_{HB} - \lambda_{LB})(\theta_{HB} - \theta_{HG})$, the greater is the scope for absorbing changes in skill supplies by changes in product mix, and the smaller the need for changes in skill intensity within sectors. The lower are σ, ϵ_{BG} and δ_{BG}, the more does a rise in the relative supply of skilled workers depress their relative wage. The effect of varying $(\lambda_{HB} - \lambda_{LB})(\theta_{HB} - \theta_{HG})$ depends on $(\epsilon_{BG}\delta_{BG} - \sigma)$, but if there is more substitutability in consumption than in production, a smaller difference in skill intensity has an effect similar to that of lower substitution elasticities.

---

[6] Apart from the introduction of δ, equation (9) is the reciprocal of the 'economy-wide elasticity of substitution between factors' in Jones (1965, eq. 15), but with the simplifying assumption of a single σ across sectors, which is common in the literature (e.g., Katz and Autor (1999) and Morrow and Trefler (2017)). With sector-specific σ's, the first term in the denominator of (9) (and the corresponding term in (2)) would be $\sum_{j=1}^{n}\left[\lambda_{Hj}(1 - \theta_{Hj}) + \lambda_{Lj}\theta_{Hj}\right]\sigma_j$. But with a common σ, the term simplifies: $\sum_{j=1}^{n}\left[\lambda_{Hj}(1 - \theta_{Hj}) + \lambda_{Lj}\theta_{Hj}\right]\sigma = \left[1 - \sum_{j=2}^{n}(\lambda_{Hj} - \lambda_{Lj})(\theta_{Hj} - \theta_{H1})\right]\sigma$.

In practice, because data on goods are always aggregated, the aggregate elasticity of substitution in production (the first term in the denominator of the rhs ratio) encompasses wage-induced changes in intra-sectoral product mix as well as in techniques of production (Davis and Weinstein, 2001; Schott, 2003). Within each sector is a relationship qualitatively similar to that across sectors described by equation (9), including the effect of changes in the relative prices of products of differing skill intensity on their relative sales and hence their relative weight in the output of the sector.

(d) Openness to trade, intermediate inputs, and mobile capital

The 2×2 framework can straightforwardly be extended to include n goods, indexed by j and with good 1 as the numeraire, as in equation (10)[7]

(10)

$$\widehat{w}_{H/L} = -\frac{1}{\left[1 - \sum_{j=2}^n \left(\lambda_{Hj} - \lambda_{Lj}\right)\left(\theta_{Hj} - \theta_{H1}\right)\right]\sigma + \sum_{j=2}^n \left(\lambda_{Hj} - \lambda_{Lj}\right)\left(\theta_{Hj} - \theta_{H1}\right)\epsilon_{j1}\delta_{j1}}\widehat{v}_{H/L}$$

whose denominator is again a weighted average of the common σ and of the $\epsilon\delta$ terms, with weights dependent on the variance of skill intensity across sectors. Moreover, as in other models of this sort (e.g. Arkolakis et al., 2012, pp. 115-6), the n output sectors can include both final goods and intermediate goods, with the demand system being reinterpreted to cover both consumers and firms.

Introducing intermediates requires other extensions of the model, for brevity described only in words here. As for example in Eaton and Kortum (2002), the output of each good in each country is produced with labour and imported intermediate inputs, whose cost to the importing country varies with its geographical location and other costs of trade, but is unlikely to be correlated with its relative supplies of skilled and unskilled labour. For the purpose of calculating the price-ratio elasticity, labour input in each sector should in principle include the indirect contribution of domestically-sourced non-traded intermediate inputs, with the labour cost of domestically-sourced traded intermediate inputs (import substitutes or exportables) as well as that of imports being added to the other-cost wedge.[8]

---

[7]In a multi-country setting and with γ's and β's varying by destination country, the overall elasticity of relative sales with respect to relative purchaser prices, ϵ_{j1}, is a weighted average of the bilateral elasticities (the approximated ϵ_{BG} in equation (5), without country superscripts), where the weights are output shares averaged across the two sectors j and 1 - see section V for further details.

[8]Since domestic and imported inputs are usually imperfect substitutes, distinguishing non-traded from traded intermediate inputs would require an arbitrary cut-off value of the elasticity of substitution.

The definitions of some variables thus change slightly: for example c becomes direct plus non-traded indirect labour cost, with θ referring to its skilled share. Production also requires capital, which is assumed to be internationally mobile (for reasons explained in Wood (1994, section 2.2) and supported by the evidence in Caselli and Feyrer (2007)), and thus plays a role similar to that of traded intermediate inputs.

Openness to international markets affects the elasticity of relative wages with respect to relative skill supplies through the values of ϵ_{j1} and δ_{j1}. Both these parameters are increased by lower trade costs, though in different ways. What matters for ϵ_{j1}, by affecting home market shares and access to export markets, are international trade barriers of all types, whether or not proportional to ULCs. What matters for δ_{j1} is only trade costs that are not proportional to ULCs, but including internal as well as international trade costs (plus the costs of traded intermediates and mobile factors). (A closed economy would be a special case of (10) in which, because $s_{BG} = 1$ in (5), $\epsilon = \gamma$, but internal per-unit OCWs still made $\delta \leq 1$.)

Trade in intermediate inputs creates additional links between openness and the wage-skill supply elasticity. It reduces δ_{j1} and thus tends to lower (i.e. make more negative) the wage-skill supply elasticity because traded intermediates are a component of the relative costs of goods that is independent of labour costs – the same being true of capital. Cutting the other way, trade in intermediates increases ϵ by raising measured β (relative to γ). From an 'Armington' perspective, availability of the same inputs to countries makes their outputs closer substitutes (e.g. the failings of a country's zip industry need no longer make its jeans less desirable than those of some other country). From an 'Eaton-Kortum' perspective, importing intermediates makes countries even more specialised in activities in which their skill endowments give them a comparative advantage, increasing the cross-country variability of relative production costs.[9]

In summary, greater interaction with the world economy unambiguously raises the wage-skill supply elasticity by increasing ϵ_{j1} – via lower home-firm shares of the domestic market and higher shares of exports in home-firm sales, and via higher measured β (in the ways explained in the previous paragraph). It has an ambiguous effect, however, via δ_{j1}, whose value is increased by lower (per-unit) trade costs but reduced by greater use of imported intermediate inputs and internationally mobile capital.

---

[9]The effect is similar to that of a lower value of the θ parameter in Eaton and Kortum (2002). Greater use of imported intermediates may be stimulated by reduction either of trade costs or of what Anderson et al. (2006) call 'cooperation costs' and Baldwin and Robert-Nicoud (2014) call 'coordination costs'.

(e) Other influences on the wage-skill supply elasticity

Equation (10) conveys the main point of this paper, which is that the elasticity of relative wages with respect to relative skill supplies depends on the degree of openness of the economy concerned. Relative wages are of course subject to many other influences, and of special interest in this paper are two mechanisms by which variation in skill supplies can affect wages in ways other than the substitution processes embedded in equation (10) and hence alter the relative wage-skill supply elasticity.

One is directed technical change, as in Acemoglu (2007), who argues that increases in the supply of a factor stimulate innovation to use it more productively. Denoting factor-augmenting technical change by $\widehat{r}_{H/L}$ (the rate of improvement for skilled labour minus that for unskilled labour), its net effect on the relative demand for skilled labour is $-(1-\widetilde{\sigma})\widehat{r}_{H/L}$, where $\widetilde{\sigma} \equiv \left[1 - \sum_{j=2}^{n}(\lambda_{Hj} - \lambda_{Lj})(\theta_{Hj} - \theta_{H1})\right]\sigma$, and is positive if $\widetilde{\sigma} > 1$ (which requires $\sigma > 1$).[10] In a framework with a single final output sector, Acemoglu derives the elasticity of $\widehat{r}_{H/L}$ with respect to $\widehat{v}_{H/L}$ as being equal to the elasticity of substitution between skilled and unskilled labour minus one.[11] In our multi-sector framework, the bias of directed technical change depends on the economy-wide elasticity of substitution between factors (note n. 6), which is the full denominator of equation (10) and can be written compactly as $\phi_{H/L}$. Combining these elements yields the middle term in the numerator of equation (11)[12]

$$(11) \quad \widehat{w}_{H/L} = -\frac{1 + (1-\widetilde{\sigma})\left(\phi_{H/L} - 1\right) - \left[\widetilde{\omega} + \sum_{j=2}^{n}(\lambda_{Hj} - \lambda_{Lj})(\mu_j - \mu_1)\right]\pi_v}{\phi_{H/L}}\widehat{v}_{H/L}$$

The other mechanism, established by Caron et al. (2014), is that the relative demand for skill-intensive goods is income-elastic, which, in combination with the effects of increases in skill supply on income, implies that skill supply increases will be partly absorbed by induced shifts in demand. The final term in the numerator of equation (11) embodies this mechanism, with π_v being the elasticity of per capita income with respect to relative skill

---

[10] A positive rate of skill-augmenting technical change, for example, pulls the relative wage of skilled workers in two opposite directions: downwards, by increasing their effective relative supply, and upwards, by making them relatively more efficient and so increasing the demand for them. If the elasticity of substitution exceeds unity, the relative supply boost is outweighed by the relative demand boost.

[11] Used also by Morrow and Trefler (2017, eq. 25). Intuitively, a higher elasticity increases the incentive for innovation to expand the effective supply of a factor by diminishing the consequent reduction in its marginal product. The result is from eq. (22) of Acemoglu (2007), but assuming linear R&D technology.

[12] This mechanism, too, might interact with openness: Acemoglu et al. (2015) analyse how offshoring alters the skill bias of innovation.

supply, $(\mu_j - \mu_1)$ the difference between the income elasticity of demand for good j and that for the reference good 1, and $\widetilde{\omega}$ the 'aggregate' effect of income on the skill intensity of the composition of demand for varieties of each good (needed because the data on goods are not detailed enough to include every variety)[13] – section A.IV of the Online Appendix contains a fuller discussion. The combination of directed technical change and income elasticity, as equation (11) shows, might thus cause the wage-skill supply relationship to be much more elastic (i.e. less strongly negative) than would be expected on the basis only of the EHO reallocation mechanisms.

(f) Other influences on relative wages

Other forces that influence the relative wages of skilled and unskilled workers in open economies can be thought of as extra terms on the rhs of (11). One such influence of special interest in any HO model is variation in world prices and trade costs (holding skill supplies constant), which can be described by

$$(12) \quad \widehat{w}_{H/L} = \frac{1}{\phi_{H/L}} \left[\sum_{j=2}^{n} (\lambda_{Hj} - \lambda_{Lj}) (1 - s_{j1}) (\beta_{j1} - \gamma_{j1}) \left(\widehat{p}^*_{j/1} + \widehat{T}_{j/1} \right) \right]$$

where p_j^* is the world price of good j and $T_j = \frac{p_j}{p_j^*}$ is the trade cost ratio (greater than unity if j is an import substitute and less than unity if j is an export good). The term $1/\phi_{H/L}$ is the absolute value of the elasticity defined by equation (10), which translates changes in the relative demand for skilled and unskilled workers into change in their relative wages. The terms to the right of $1/\phi_{H/L}$ show how changes in the local prices of foreign goods alter the relative demand for skilled and unskilled workers. The expression $(1 - s_{j1}) (\beta_{j1} - \gamma_{j1})$ describes the effects of changes in the relative price of foreign varieties of goods j and 1 on the relative sales of domestic varieties of these goods, which $(\lambda_{Hj} - \lambda_{Lj})$ translates into changes in relative labour demands.[14]

Though in the same direction as in HOS, the effect of foreign prices on wages is smaller

---

[13]$\widetilde{\omega} = \left[1 - \sum_{j=2}^{n} (\lambda_{Hj} - \lambda_{Lj}) (\theta_{Hj} - \theta_{H1}) \right] \omega$; with ω being assumed common for all goods.

[14]The s_{j1}, $\widehat{p}_{j/1}$, $\widehat{T}_{j/1}$, are weighted averages across markets. The effect of $\widehat{p}_{j/1}$ and $\widehat{T}_{j/1}$ on relative wages increases with β_{j1}, the degree of substitutability among varieties (though not so fast as equation (12) may suggest, since a rise in β_{j1} also makes $\phi_{H/L}$ smaller). The effect of $\widehat{p}_{j/1}$ and $\widehat{T}_{j/1}$ on relative outputs decreases with the market share of home producers: with a larger s_{j1}, rises (say) in the prices of foreign varieties of good j cause less of an increase in the relative sales of j by home producers because the scope for substitution towards domestic varieties is reduced by the initially smaller sales of foreign producers. A fuller exposition of equation (12) is in sections 3.4 and 5.3 of Wood (2012).

and less direct, because the local prices of foreign varieties, being imperfect substitutes for domestic varieties, do not directly influence their prices.[15] If domestic and foreign varieties were perfect substitutes ($\beta = \infty$), equation (12) would reduce to HOS (though with the degree of amplification of relative wages with respect to relative foreign prices increased if $\delta < 1$), and relative wages would be independent of relative skill supplies.

Lack of data in WIOD on world prices or (a full enough set of) trade costs precludes estimation of equation (12) in this paper, whose focus is on the elasticity of relative wages with respect to skill supplies described by equations (10) and (11). In estimating (10), however, it will be necessary to recognise the existence of (12). Greater openness increases the elasticity in (10) (i.e., makes it less negative), but it also changes the $T_{j/1}$ in (12), in directions likely to be related to skill supplies, creating another sort of relationship between wages, skill supplies and openness that could bias the estimation of (10).

In estimating (10), it will also be important to recognise the existence of many other possible influences on relative wages that are likely to cause this equation not to fit the data well. Particular attention has been given in the literature to skill-biased technical change and labour market institutions. As with world prices and trade costs, moreover, some of these influences depend on openness, including economies of scale (Epifani and Gancia, 2008), offshoring (Grossman and Rossi-Hansberg, 2008), firm heterogeneity (Harrison et al., 2011; Burstein and Vogel, 2016), and the transfer of skill-intensive technology or products to skill-scarce countries (Wood, 2002; Burstein et al., 2013).

III Empirical strategy and data

The empirical analysis in the rest of this paper has two objectives. The first is to assess whether and to what extent the relative wage-skill supply elasticity varies among countries and periods in the ways suggested by our theory. The second is to assess whether the set of determinants of this elasticity discussed in the theory section provides a good explanation of the size of the observed average elasticities in Figure 1.

Sections IV and V address the first objective. They examine how variation in openness and in price-ratio elasticities across countries and over time is related to variation in the strength of the inverse relationship between relative wages and relative skill supplies, using

---

[15]If $\beta_{BG} = \gamma_{BG}$, including the Cobb-Douglas case of both parameters being unity (Abrego and Whalley, 2000), changes in foreign prices would not affect wages, and if $\beta_{BG} < \gamma_{BG}$ the direction of the effect would be reversed. β_{BG} and γ_{BG} act in opposite directions because (for example) a rise in the price of foreign varieties both raises the demand for domestic varieties and reduces demand for all varieties of the good in total.

multi-country panel data. Section IV explores reduced-form relationships, as a check on the prima facie consistency of the data with the theory. Section V applies a more stringent structural estimation approach to isolate and quantify the openness and price-ratio elasticity mechanisms. Section VI moves on to the second objective, bringing in the directed technical change and income elasticity of demand mechanisms. It deduces the likely size of the effect of these two mechanisms from the gap between the reduced-form and structural estimates in sections IV and V, and extends the structural estimation to assess whether in combination their impact is plausibly of the size suggested by this gap.

Our data are from the World Input-Output Database (WIOD), described by its creators in Timmer (2012). The core of WIOD is annual international input-output tables for 1995-2009 and 35 industries connecting 40 countries that produce 85% of world GDP, plus a composite rest of the world. Trade flows are computed from these tables. Values are at basic prices, with information on trade and transport margins for internal and international transactions. Price deflators can be computed from the same tables valued at previous year's prices.

WIOD's auxiliary SEAs provide information on employment and wages for three skill categories of workers in each country, industry and year. Skill is measured by schooling, using the International Standard Classification of Education (ISCED). 'Low skilled' workers are ISCED categories 0, 1 and 2 (everything below completed upper secondary). 'Medium skilled' are ISCED categories 3 and 4 (complete upper secondary and some tertiary, but below a college degree), and 'high skilled' are categories 5 and 6 (a 2-4 year college degree, or its vocational equivalent, and above). As in Timmer et al. (2014) and Morrow and Trefler (2017), we reduce the three WIOD skill categories to two by combining 'low-skilled' and 'medium-skilled' into 'unskilled' (those with less than a college degree).

WIOD wage and employment data were assembled from national labour force surveys and censuses, not previously collated in this form. We derive wage rates for each skill level, country and year as labour compensation divided by hours worked.[16]

An important maintained assumption of our analysis is that WIOD data on economy-wide employment by skill category are a good proxy for skill supplies (or endowments). Because of differences in participation rates, unemployment rates and working hours (the unit of employment in WIOD), there are differences across countries and periods in the relationship between total employment at specified levels of education and the numbers of adults with those levels of education (which is arguably the best measure of endowments). This raises

_____

[16]Though these data are by far the best available, their accuracy is open to doubt, especially in poorer countries with large numbers of self-employed workers, for whom wages comparable to those of employees had to be imputed. In some countries, gaps had to be filled by using data from other similar countries. In addition, of course, the quality of schooling varies widely.

identification issues, since observed variation in relative employment of skilled and unskilled workers arises not only from supply shocks but also from demand shocks and institutional influences on relative wages. However, the observed strong correlation with variation in the corresponding population ratios in Barro and Lee (2013) suggests that the variation in skilled employment ratios in WIOD is mainly due to variation in skill supplies.[17]

In the reduced-form analysis we add controls for average income, additional factors of production, and labour market institutions. Table A.1 in the Online Appendix reports summary statistics and Table A.2 is a correlation matrix for the main variables used in the regressions.

IV Reduced-form relationships

Our EHO model makes two novel testable predictions about variation among countries and over time in the wage-skill supply elasticity:

1. This elasticity is reduced (brought closer to zero) by more openness to trade, which increases the influence of substitutability among varieties relative to substitutability among goods and thus reduces the size of the relative price changes needed to achieve skill-supply absorbing changes in output mix.

2. This elasticity is reduced also by a higher price-ratio elasticity – the responsiveness of relative purchaser prices of goods to relative labour costs – which permits the relative price changes needed for skill supply-absorbing changes in output mix to be achieved with smaller changes in relative wages.

To assess the first prediction in reduced form, we use the commonest measure of openness: the ratio of a country's total trade (exports plus imports) to its total output, which we adjust for country size by using the residuals of a cross-section regression on population (both in logs) for each year in the sample. This openness measure is clearly related to the theoretical construct summarised in equation (5), since it is higher import and export shares that give substitutability among varieties more influence on the elasticity of demand.

The country size adjustment is also theory-based. Trade ratios are higher in smaller countries mainly because they have more non-competing imports – producing in fewer sectors

---

[17]The correlation between the Barro-Lee adult population measure of relative skill supply and the WIOD employment measure of relative skill supply is about 0.7, but it depends mainly on cross-country variation. Replacing the (annual) employment measure with the Barro-Lee (five-yearly) measure in reduced-form panel regressions yields weak results.

as a result of less scope for realising economies of scale and less diverse natural resources. The elasticity of demand in our model depends on the share of domestic producers in markets in which they compete: these shares rise with the height of barriers faced by foreign suppliers, which may be no lower in less populous countries than in more populous ones.[18,19]

To assess the second prediction, we estimate the price-ratio elasticity from the relative size of labour costs (referred to earlier as unit labour costs, or ULCs) and costs that do not vary in proportion to wages (the other-cost wedges, or OCWs). The higher the ratio of ULCs to OCWs, the more responsive are relative prices to changes in relative wages (though the price-ratio elasticity also depends on the degree, η, to which the relative OCWs of goods vary with their ULCs, on which we lack information). In relation to equation (8), our measure of this simplified price-ratio elasticity, $\widetilde{\delta}$, is thus $1/(1+\tau)$, with $\tau = {}^{OCWs}/_{ULCs}$, and is calculated for the economy as a whole.

A key issue is how to measure τ at country level. We confidently put internal trade costs and net taxes into the OCWs, since they drive a wedge between ULCs and purchaser prices, whatever the location of the purchaser. For each country and year we sum across sectors internal transport margins and net product taxes in WIOD's national Supply-and-Use tables. Given the assumption in this paper that capital is internationally mobile, we also assign WIOD's "capital compensation" to the OCWs, since it too creates a wedge between ULCs and purchaser prices regardless of where products are sold.

To calculate τ for each country and year, we then divide the sum of these OCW elements by the total national wage bill (our proxy for economy-wide ULCs). In principle, the OCWs should also include foreign trade costs and traded intermediate inputs, but the only information in WIOD on foreign trade costs is the international transport margin, and there was no practical way of isolating traded intermediates (which include domestically supplied inputs that are import-competing or exportable). In the reduced-form analysis, the influence of the omitted per-unit foreign trade costs and traded intermediates on the wage-skill supply elasticity should be at least partly picked up by our measure of openness.

The general form of our reduced-form regressions is:

$$(13) \quad \ln(w) = \alpha + \beta_1 \ln(v) + \beta_2 \ln(o) + \beta_3 \left(\ln(v) \times \ln(o)\right) + \beta_4 \widetilde{\delta} + \beta_5 \left(\ln(v) \times \widetilde{\delta}\right) + \varepsilon$$

---

[18]Because population size is correlated with land area, longer internal distances could create higher barriers in bigger countries, but not necessarily (Ramondo et al., 2016). Gravity-based measures of trade barriers control for internal distance (e.g., Anderson and Yotov (2010)).

[19]Though related in these ways to the theory, as well as having the advantage of being familiar from other empirical studies, our trade-to-output measure of openness does not correspond exactly with the specification of the equations in section II. However, we show in section V that using a measure of openness derived directly from the equations yields qualitatively similar reduced-form results.

where w is the relative wage of skilled workers, v is their relative supply, and o is openness. Our two EHO predictions suggest positive signs on the coefficients β_3 and β_5 on the v-o and v-$\widetilde{\delta}$ interactions, since higher o and $\widetilde{\delta}$ offset the negative wage-skill supply elasticity. The variables are transformed into deviations from their means, so that the coefficient on each variable alone (β_1, β_2, and β_4) measures its effect evaluated at the sample mean of the other interacted variable – e.g., β_2 is the wage-openness elasticity evaluated at the average skill supply ratio.

A limitation of this specification is that a positive β_3 could pick up more than the effect of greater openness on the wage-skill supply elasticity, as explained in section II. It could also pick up the effect of greater openness on relative goods prices, which tends to increase the relative wage of skilled workers in a high-v country, but to decrease it in a low-v country. Formally, in the EHO model, the direction of the changes in T's in equation (12) is likely to depend on v, but without information on the T's for specific goods we cannot disentangle this effect from the also-positive effect of greater openness on the wage-skill supply elasticity.

Nor does this specification identify the directed technical change and income elasticity of demand mechanisms. Both these mechanisms are deeply embedded and hard to disentangle from the EHO reallocation channels in reduced form (though we estimate them structurally in section VI). Directed technical change in the numerator of equation (11) depends on the same parameters as the reallocation mechanisms in the denominator. As regards the income elasticity of demand mechanism, there are likewise no obvious causes of variation in $\widetilde{\omega}$ and μ_j among countries and periods, and too many possible causes of variation in π_v (the effect of a greater relative supply of skill on per capita income), including openness to trade.

These theoretical interconnections and the limitations of our data preclude accurate reduced-form identification of our openness and price-ratio elasticity mechanisms. But to make a broad assessment of whether and to what extent their influence leaps out from the data, we use the general regression framework in (13) in three different ways: across countries in levels (with variables averaged over time for each country, as in Figure 1(a)); across countries in full-period changes, as in Figure 1(b) (with $\ln(o)$ and $\widetilde{\delta}$ at their initial 1995 values); and as a panel with annual data, including country fixed effects and year dummies. The panel approach reduces bias from unobserved country-specific and time-invariant influences and global trends, and increases the sample size from 40 to 600, making it possible to control for other potentially confounding influences.

Table 1 presents the results. For each type of variation in relative wages and skill supplies, we report estimates of the wage-skill supply elasticity alone (columns (1), (3), and (5)), as well as the specification in (13) with the v-o and v-$\widetilde{\delta}$ interaction terms (columns (2), (4),

20

and (6)). For the panel specification only, we also report results including six other control variables, both on their own (column (7)) and interacted with v (column (8)).

The coefficient on v is negative in all specifications, confirming the inverse relationship between relative wages and skill supplies in Figure 1.[20] Its size and significance vary across specifications, but the significant values are never far from those in Figure 1 (between -0.17 and -0.35). The coefficient is small and insignificant in the first panel regression (column (5)), which includes only the fixed effects and year dummies, but becomes much larger and significant with the controls for openness and the price-ratio elasticity in column (6).

The coefficients on the openness variable alone, which measure its relationship with w at the sample average of v, vary widely among specifications: the significant but opposite-signed values in the small-sample cross-country and long-change results are more likely to reflect the influence of omitted variables than to shed light on the accuracy of alternative theories of trade and wages.

The v-o interaction terms support the EHO theoretical predictions. The estimated interaction coefficients are consistently positive, suggesting that higher trade openness weakens the otherwise negative wage-skill supply elasticity. They are mostly about 0.4, though less precisely estimated across countries and in long changes.[21]

---

[20]The sign and size of the wage-skill supply elasticity in the full-period changes specification are confirmed when using exogenous net migration flows as an instrument for changes in skill supplies. Figure A.1 in the Online Appendix portrays the results and the notes provide details on the identification strategy.

[21]Standard errors are clustered at the country level and bootstrapped because the openness variable is a "generated regressor" (the residuals of cross-section regressions of the trade-output ratio on population).

Table 1: Wage-skill supply elasticity and EHO in reduced form

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Cross-country		14-year changes		Panel			
$\ln(v)$	-0.269***	-0.275***	-0.260*	-0.352**	-0.0781	-0.177**	-0.178**	-0.117
	(0.0751)	(0.0704)	(0.146)	(0.155)	(0.0976)	(0.0744)	(0.0788)	(0.0912)
$\ln(o)$		-0.294**		0.0163***		-0.0461	-0.0200	-0.0469
		(0.150)		(0.00543)		(0.0645)	(0.0661)	(0.0546)
$\ln(v) \times \ln(o)$		0.386		0.825		0.391***	0.396***	0.376***
		(0.334)		(0.518)		(0.103)	(0.105)	(0.0748)
$\widetilde{\delta}$		-0.879		-1.70e-05		-0.587**	-0.689*	-0.472
		(0.598)		(0.0202)		(0.256)	(0.413)	(0.298)
$\ln(v) \times \widetilde{\delta}$		1.671*		-0.987		0.745**	0.745***	0.790*
		(0.935)		(1.748)		(0.301)	(0.263)	(0.428)
Obs	40	40	37	37	600	600	600	600
R^2	0.355	0.535	0.106	0.373	0.069	0.332	0.353	0.408

Notes: All variables are within-country averages in columns (1) and (2). In columns (3) and (4), $\ln(w)$ and $\ln(v)$ are in 2009-1995 differences, whereas values for $\ln(o)$ and $\widetilde{\delta}$ are for the year 1995. Country fixed effects and year dummies are included in columns (5) to (8). Column (7) includes control variables (see text and Table A.1), and column (8) further adds interactions of each control variable with $\ln(v)$. Heteroskedasticity-robust standard errors are reported in parenthesis in columns (1) and (3), standard errors bootstrapped with 500 replications are in parenthesis in columns (2) and (4), and clustered-bootstrapped standard errors are in parenthesis in columns (5) to (8). Significant at: *10%, **5%, ***1% level.

The panel estimates in column (6) are used in Figure 2 to show graphically how the wage-skill supply elasticity varies with trade openness, holding the level of $\widetilde{\delta}$ at its average value. In a low-trade economy (such as Brazil: average $\ln(o) = -0.51$), a 10 percent increase in relative skill supply is associated with a 3.8 percent drop in the skill premium, and in a hypothetical almost-closed economy with a striking (but very inaccurately estimated) drop of 20 percent. By contrast, countries in the upper quartile of openness (such as Germany: average $\ln(o) = 0.25$) exhibit relative wage insensitivity – no significant variation in the skill premium with skill supplies. The cross-country and long-change regressions (columns (2) and (4)) suggest even greater variation of the elasticity with trade openness: their smaller samples reduce precision without altering the economic implications. Using the long-change results in column (4), a 10 percent increase in the relative skill supply is associated with a statistically significant 8 percent drop in the skill premium for Brazil ($\ln(o) = -0.59$ in 1995), but no significant change for countries in the upper tercile of openness.

Figure 2: Wage-skill supply elasticity and openness at average $\widetilde{\delta}$
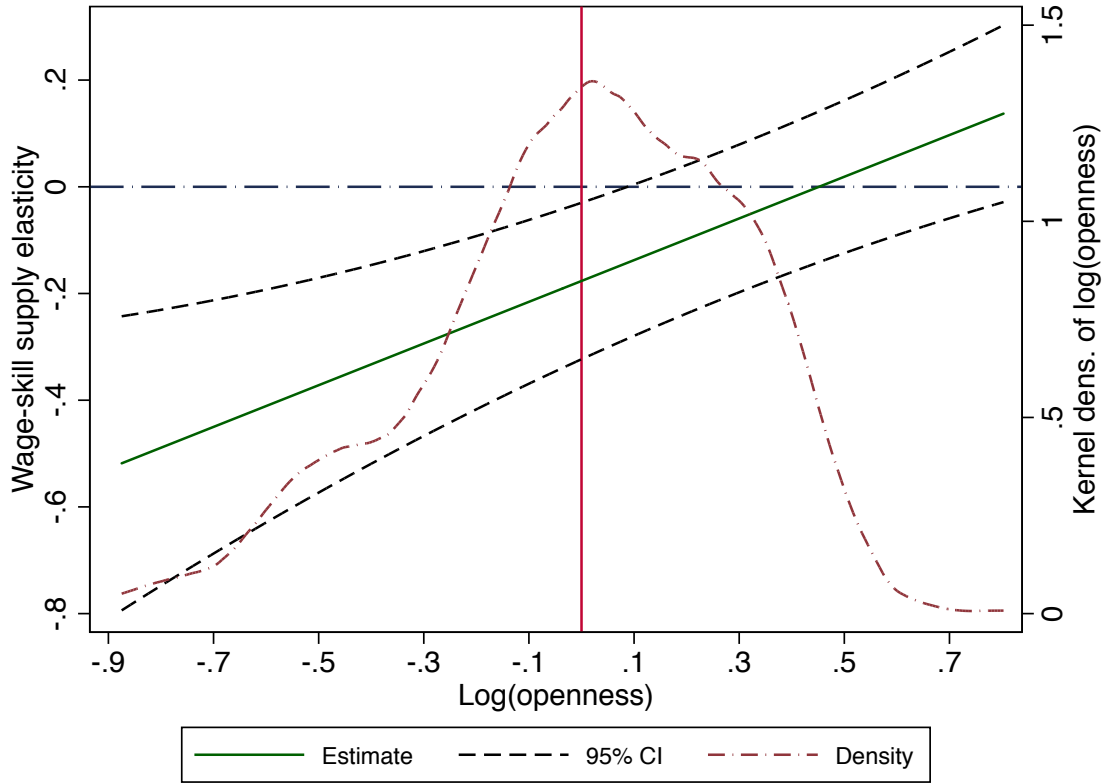
These elasticity-reducing effects of greater openness are not materially altered if the trade ratio is calculated on a value-added rather than gross basis. Nor, as shown in Table A.3 of the Online Appendix, are they much altered if the numerator excludes trade in intermediates (whose effect on the wage-skill supply elasticity in theory is ambiguous), trade with China (which has been found to have strong effects on local wages and labour markets (Autor et al., 2016)), or trade in equipment proxying for transfers of technology (Burstein et al., 2013).

The coefficients on the price-ratio elasticity alone, which measure its relationship with w at the sample average of $\ln(v)$, are mainly negative (though zero in the long-changes specification), perhaps because of capital-skill complementarity – since capital is assumed to be mobile and payments to capital reduce $\widetilde{\delta}$.

As predicted by our EHO model, the v-$\widetilde{\delta}$ interaction coefficients are mainly positive and significant, but are twice as large in the cross-country as in the panel specification and insignificantly negative in the long-change specification. These differences may partly reflect the omission of international trade costs and traded intermediates, whose effect (and relationship with the openness variable) could differ across specifications.

As well as being less robust, the role of variation in price-ratio elasticities is smaller than that of trade openness. This can be seen in Figure A.2 in the Online Appendix, where we trace out how the wage-skill supply elasticity varies with our measure of $\widetilde{\delta}$, using the panel estimates in column (6) and holding log openness at its average value. The wage-skill supply elasticity rises more slowly with the price-ratio elasticity than with openness. Even so, if for example Turkey had average openness and increased its price-ratio elasticity from the low actual average value of 0.35 (10th percentile) to 0.55 (90th percentile), the relative wages of its skilled workers would no longer be sensitive to changes in skill supplies.

Other observable determinants of relative wages may be correlated with skill supplies, openness or price-ratio elasticities, and could thus bias our estimates of the v-o and v-$\widetilde{\delta}$ interaction coefficients. To check this, we add to our panel regression six more variables (defined more precisely in the notes to Table A.1): the share of low-skilled workers in the unskilled (low + medium) aggregate, and the share with no education in the low category (computed from Barro and Lee (2013)); the ratio of a country's land area to its unskilled labour supply; the ratio of the capital rental rate to the average wage; GDP per capita; and an EU membership dummy (because EU countries are over-represented in our sample).[22]

As can be seen from columns (7) and (8) of Table 1 (the difference between which is that the latter includes also the interactions of the six control variables with v), the addition of these other determinants of relative wages has no substantial effect on the sizes of the v-o and v-$\widetilde{\delta}$ interaction coefficients (the coefficients on the extra variables in the column (7) regression are in Table A.4 in the Online Appendix). The small samples in the cross-country and long-change regressions made the addition of more variables less informative, but neither individually nor together did the six extra variables affect substantially the interaction coefficients (columns (1) to (4) of Table A.4).

In summary, our reduced-form evidence is consistent with the predictions of our EHO model concerning the wage-skill supply relationship. Both greater trade openness and higher price-ratio elasticities attenuate the negative elasticity of the skilled wage premium with respect to skill supplies – and with high openness and high price-ratio elasticities wages vary only slightly or even not at all with variation in skill supplies.

---

[22]We experimented also with measures of labour market rigidity, including unionisation, the Employment Protection Legislation index from the OECD (2013) and the "Labour Freedom' index from the Heritage Foundation, but all of them substantially reduced our sample size in non-random ways. For example, the differences between columns (6) and (7) in Table A.3 arise mainly from the change in sample rather than from adding the unionisation variable.

V   Structural relationships

To isolate and quantify the openness and price-ratio elasticity mechanisms of the EHO model, we now estimate them directly. We assess whether the variation in relative wages with respect to skill supplies that they predict (the right-hand side of equation (10)) matches the observed variation in the data. We continue to abstract from the directed technical change and income elasticity of demand mechanisms in the numerator of (11), which the structural estimation is extended to include in section VI.

Our focus is on the second term of the denominator of equation (10) (abbreviated in (11)), which contains the purchaser-price and price-ratio elasticities parameters and governs the skill-supply-absorbing response of sectoral output mix. To the extent that inter-sectoral output adjustment falls short of absorbing the whole of any change in skill supply (as it would in a one-cone HOS model), the first – entirely standard – term in the denominator fills the gap through intra-sectoral adjustments of technique or detailed output mix.

(a) Estimation of parameters

A first step is to measure the parameters in equation (10). The labour use and cost shares (λ's and θ's) for skilled and unskilled labour can be calculated from the WIOD SEA for each country, sector and year. We use direct labour inputs, because of the difficulty mentioned earlier of separating traded from non-traded intermediate inputs (only the latter would be relevant for our purpose).

To measure the price-ratio elasticity δ_{j1}, we assume for simplicity that all OCW's are per-unit and hence set $\eta_{j1} = 0$. On this assumption, the price-ratio elasticity – of relative purchaser prices with respect to relative ULC's – is (as on p.10 of Wood (2012), but omitting the sector-pair subscripts for clarity)

$$(14) \quad \delta = \frac{1 + \frac{\tau}{2}\left(\sqrt{\frac{c}{t}} + \sqrt{\frac{t}{c}}\right)}{1 + \tau\left(\sqrt{\frac{c}{t}} + \sqrt{\frac{t}{c}}\right) + \tau^2}$$

where $c = \frac{c_j}{c_1}$ and $t = \frac{t_j}{t_1}$ are relative ULC's and OCW's, respectively; and $\tau = \sqrt{\tau_j \tau_1}$, where $\tau_j = \frac{t_j}{c_j}$. If $t = c$, this expression conveniently reduces to $\widetilde{\tau} = \frac{1}{1+\tau}$, which was how we approximated country-level price-ratio elasticities in the reduced-form analysis. Here we follow (14), estimating at sector level, but still including in t only internal trade costs, (net) taxes and capital compensation.

For the common elasticity of substitution in production, σ, we impose the value of 1.67 estimated by Morrow and Trefler (2017) from the same WIOD data, using a simplified version of Acemoglu (2007)'s equation for skill-biased technical change (see footnote 11). We run sensitivity checks by varying σ within the range of 1 to 3 that Katz and Autor (1999) consider plausible.[23]

To measure ϵ, we first estimate the demand parameters γ and β, applying the approach pioneered by Feenstra (1994). A variety is defined as good j (a WIOD sector) sold by country z to country $\check{z}$, where $z=\check{z}$ for a domestic variety, for which the WIOD international input-output table contains data on final demand flows. Following Patel et al. (2014), we proxy prices with sectoral price deflators.[24] The CES demand system is estimated with the LIML estimator and constrained search algorithm introduced by Soderbery (2015) to ensure that $\beta > 1$. Mean and median values across countries for β in each sector are reported in Figure A.3 in the Online Appendix: the overall median of the estimated β's is 2.57 (and mean 5.87), which is in line with earlier estimates.[25] The γ for each country is estimated similarly, exploiting variation across sectors in final consumption purchases and prices (see also Patel et al. (2014)): the estimates range from 1.21 to 2.21 (except for outlier Slovakia's 5.87), and, consistently with our assumptions, are generally lower than the β's.[26]

The estimated β's and γ's are then combined into the purchaser-price elasticity. With superscript z denoting the country of origin and $\check{z}$ the country of destination, the elasticity of country z's relative sales of goods j and 1 in market $\check{z}$ with respect to its relative prices in that market is, in a two-level CES demand system, a weighted harmonic mean of γ and β (Sato, 1967)

$$(15) \quad \epsilon_{j1}^{z\check{z}} = \frac{\frac{1}{s_j^{z\check{z}} s_j^{\check{z}}} + \frac{1}{s_1^{z\check{z}} s_1^{\check{z}}}}{\frac{1}{\beta_j^{\check{z}}}\left(\frac{1}{s_j^{z\check{z}} s_j^{\check{z}}} - \frac{1}{s_j^{\check{z}}}\right) + \frac{1}{\beta_1^{\check{z}}}\left(\frac{1}{s_1^{z\check{z}} s_1^{\check{z}}} - \frac{1}{s_1^{\check{z}}}\right) + \frac{1}{\gamma^{\check{z}}}\left(\frac{1}{s_j^{\check{z}}} + \frac{1}{s_1^{\check{z}}}\right)}$$

---

[23]In previous drafts of the paper, we estimated sector-specific σ's, whose average was 1.34, and whose use instead of the common value from Morrow and Trefler (2017) leaves the empirical findings largely unaffected.

[24]To alleviate measurement error, especially on the price side, we drop the top and bottom 1% of the sample in each product. To ensure a minimum level of identifying variation, we also omit purchasing-country-product combinations with data for fewer than ten seller countries throughout the sample period. We use the within-sector median β's for the missing observations. As in Patel et al. (2014), the estimated β for each country-sector is the median of 100 bootstrap replications. In estimating γ's (again as the median of 100 bootstrap estimates for each country), we omit countries with data (after dropping outliers) for fewer than ten sectors throughout the period (Bulgaria, Estonia, Indonesia, Ireland and Russia, for which we use the median estimated γ).

[25]Broda and Weinstein (2006), for instance, obtain a median estimated β (their σ) of 3.39 using a similar methodology with more detailed trade data (excluding home sales) for 73 countries.

[26]On average, 93% of sectors have β's above the γ of the country concerned (and 100% in 18 countries out of 40), though in Slovakia this share is only 20%.

where the weights depend on market shares. Specifically, $s_j^{z\check{z}}$ is the share of z's varieties in country $\check{z}$'s consumption of good j, while $s_j^{\check{z}}$ is the share of good j in country $\check{z}$'s total consumption.[27] The larger are the $s_j^{z\check{z}}$, the greater is the influence of country z's varieties on the overall prices of goods in country $\check{z}$, and hence the closer is $\epsilon_{j1}^{z\check{z}}$ to the lower $\gamma^{\check{z}}$ than to the higher β's.

The elasticity for country z producers across all markets is also a weighted harmonic mean, $\epsilon_{j1}^z = \left(\sum_{\check{z}} \frac{x_{1j}^{z\check{z}}}{\epsilon_{j1}^{z\check{z}}} \right)^{-1}$, where the weights ($x$'s) are the shares of their sales in each market. This elasticity increases with openness to trade: from the import side, easier access to foreign varieties reduces home market shares ($s_j^{\check{z}\check{z}}$), giving more weight to the generally higher β's; and from the export side, higher x's in foreign markets give more weight to the ϵ's in foreign markets, which are usually higher than in the home market. That ϵ_{j1}^z increases with trade openness is strongly confirmed in our data, both across countries and sectors and over time, which supports our interpretation of the reduced-form relationships.[28]

Estimation of equation (10) requires us to choose the reference sector 1. Since some parameters vary among sectors, this choice affects the predicted economy-wide wage-skill supply elasticity, and hence the comparisons with actual outcomes. The results reported below use as the reference sector 'Other non-metallic minerals', which has the fewest missing estimated parameters.[29] The conclusions would not be much affected by the choice of other sectors, as explained below.

Using the more strictly EHO-based purchaser and price-ratio elasticities from equations (14) and (15) in our reduced-form analysis confirms the evidence discussed in section IV. We check for this by replicating the regression in equation (13) but with the size-adjusted trade-output ratio (variable $\ln(o)$) and the simplified $\widetilde{\delta}$ being replaced by their counterparts, the purchaser-price elasticity ϵ_{j1}^z, and the price-ratio elasticity δ_{j1}, both aggregated at the country level as weighted averages across sectors j, where the weights equal the sector's share in total output. Table A.5 in the Online Appendix show the estimates, where the ϵ and the δ variables are the medians of the aggregated elasticities across reference sectors. The results are qualitatively similar to the ones shown in Table 1 – the negative wage-skill supply elasticity becomes less negative with higher ϵ and δ, with the influence of the latter

---

[27]The simplified arithmetic expression in equation (5) assumes two-level symmetrical CES utility (where the upper-level shares $s^{\check{z}}$ do not vary across goods) and two markets – domestic and foreign.

[28]Regressions of our estimated purchaser-price elasticities on sector-level trade openness deliver positive and significant coefficients both across and within country-sectors (results available upon request).

[29]We avoided choosing service sectors because more of their estimated elasticities seemed problematic and because trade data on services are less accurate.

being less robust across specifications. Correlation between the reduced-form variables and the aggregated EHO elasticities is indeed strong (coefficients of 0.44 between $\ln(o)$ and ϵ and of 0.93 between δ and $\widetilde{\delta}$), thus bolstering the theoretical interpretation of the reduced-form evidence.[30]

(b) Cross-country evidence

As in the reduced-form analysis, we perform the structural comparisons of predicted relative wages with observed ones using both cross-country and panel data. We begin with a slope test: the changes in relative wages predicted by our estimate of equation (10), given the actual changes in relative skill supplies ($\widetilde{w}$ for short), are regressed on actual changes in relative wages (w). In the cross-country analysis, 'changes' in relative wages and skill supplies refer to deviations from the relevant means across countries in each year, averaged across years for each country.

Figure 3 portrays the results of the cross-country slope tests. The left-hand panel, based on a straightforward application of equation (10), we refer to as the P-RE ('Price-Ratio Elasticity') case. The slope is almost equal to one and precisely estimated, suggesting that the model accurately explains how variation in skill supplies across countries affects the relative wages of skilled workers. The fit is far from perfect (R^2=0.38), which could reflect errors in the predictions as well as variation across countries in the determinants of actual relative wages other than skill supplies. Ten countries in the north-east quadrant, including China, have below-average skilled wage premia despite also having below-average skill abundance (implied by their above-average predicted skill premia).

The other two panels of Figure 3 use more restricted specifications of the output-mix elasticity – the second term in the denominator of (10). 'Trade Protection' (TP) retains the same specification of ϵ, but assumes all trade costs and other OCWs to be ad-valorem ($\delta = 1$), which reduces the predicted wage changes. The slope is slightly shallower and the fit almost unchanged. In the right-hand panel is the 'Armington Case' (AC), where the purchaser-price elasticity in each market is simply β (averaged across sectors), continuing to assume as in TP that $\delta = 1$. This case is equivalent to a one-level CES specification where $\gamma = \beta$, and is the standard setting in recent quantitative trade models (Costinot and Rodriguez-Clare, 2014), related gravity applications (Anderson, 2011) and output-side empirical tests of HO theory (Trefler, 1995). Its estimated slope is 40% lower than in the

---

[30]Panel regressions with country fixed effects and year dummies of $\ln(o)$ on ϵ and of $\widetilde{\delta}$ on δ deliver slope coefficients of, respectively, 0.474 (standard error=0.144), and 0.69 (standard error=0.131).

P-RE case. These comparisons suggest that variation in trade shares and, to a lesser extent, in per-unit costs are key to explaining how skill premia vary with skill abundance across countries.

Figure 3: Slope tests – Cross-country estimates
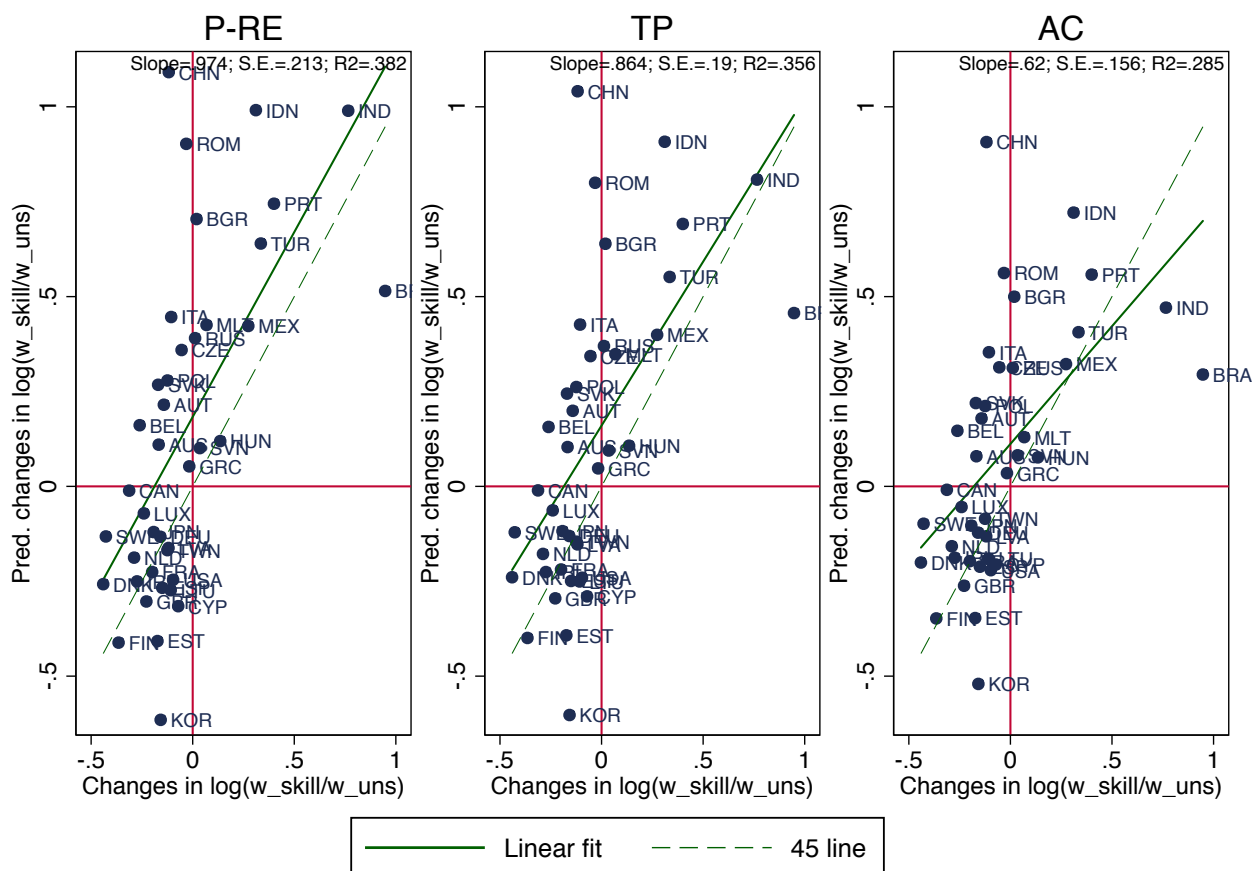
Notes: All regressions include a constant. Heteroskedasticity-robust standard errors.

Both predicted and actual changes in wages are arguably measured with error. In such a situation, Klepper and Leamer (1984) show that bounds of the true slope coefficient are given by the estimated slopes of the 'direct' ($\widetilde{w}$ on w) and 'reverse' (w on $\widetilde{w}$) regressions. The upper panel of Table 2 shows these bounds. In the P-RE case, the reverse regression yields a far smaller (0.4), though still positive and significant, slope than the direct regression. This difference is amplified by the ten countries where the skill premium is unusually low. Without them, the estimated slope is between 0.58 and 1.06, with a much better fit (R^2=0.6). The bounds of the slope parameter in the TP and AC specifications are narrower, and the 'reverse' slopes are similar in all three cases.

Another test of fit is the variance ratio (Trefler, 1995; Blum, 2010): the variance of

predicted changes in relative wages ($\widetilde{w}$) divided by the variance of actual changes (w). The lower panel of Table 2 shows that the variance of relative wages across countries predicted by our P-RE and TP specifications is twice as large as the actual variance of relative wages (though only slightly larger with the AC specification, which predicts a lower wage-skill supply elasticity). Plausible explanations for this over-prediction are the omitted directed technical change and income elasticity of demand mechanisms. Labour market and social security institutions may also narrow differences in skill premia across countries.

Table 2: Bounds and Variance-ratio tests – Cross-country

	P-RE		TP		AC	
Bounds:	$\tilde{w}$ on w	w on $\tilde{w}$	$\tilde{w}$ on w	w on $\tilde{w}$	$\tilde{w}$ on w	w on $\tilde{w}$
	0.974***	0.409***	0.864***	0.431***	0.620***	0.489***
	(0.213)	(0.106)	(0.190)	(0.117)	(0.156)	(0.144)
Obs	40	40	40	40	40	40
R^2	0.382	0.382	0.356	0.356	0.285	0.285
Variance ratio: $\frac{Var(\tilde{w})}{Var(w)}$						
	2.383		2.002		1.268	

Notes: All regressions include a constant. Heteroskedasticity-robust standard errors are in parenthesis. Significant at: *10%, **5%, ***1% level.

Our results survive a number of robustness checks.[31] To see if they are driven by specific time periods, we apply the slope test to each year in the sample between 1995 and 2009. As shown in Figure A.4 in the Online Appendix, the slope increases slightly over time in all three specifications, although not significantly. We also replicate the slope test with different choices of the reference sector. Figure A.5 in the Online Appendix reports the estimated slopes for the three model specifications. The point estimates do not vary markedly across reference sectors and they preserve the ordering across specifications, with P-RE generating a higher 'direct' slope than the two alternatives.

(c) Panel evidence

To measure changes in relative wages and skill supplies within countries over time (the 'hat' terms in equation (10)), we follow Blum (2010) and take annualised differences (in logs) over windows of 5, 10 and 14 years. The value of the wage-skill supply elasticity (the

---

[31]The variance of predicted wage changes, and hence the slope of their regression on actual wage changes, decreases mechanically with the chosen value of σ (which is 1.67 in the results reported). With the P-RE specification, raising σ from 1 to 3 reduces the slope from 1.5 to 0.5, though the R^2 remains at 0.38 (results available upon request).

ratio term on the rhs of equation (10)) is taken at the start of each window. Predicted wage changes are then regressed on actual changes and, in the 5- and 10-year regressions, also on country fixed effects and time dummies. Table 3 reports the results of the (direct and reverse) slope and variance-ratio tests.

Starting with the baseline P-RE column, the estimated slope is positive and significant in both the 5- and 10-year windows, and is greater in the longer window, though still well below unity. The same is true of the reverse regressions, whose slopes are much closer to those of the direct regressions than in the cross-country test. In the longest possible window of 14 years, the slope estimated on the full sample is near zero. When three extreme outliers are omitted, as with Figure 1(b), the slope is positive and close to those in the shorter windows, though less precisely estimated. The variance ratio is below unity and roughly constant across the three time windows, suggesting that changes in actual skill premia over time are amplified by forces omitted from equation (10).

As in the cross-country tests, but less so, the direct slopes of the AC specifications in the last four columns of Table 3 are lower than for P-RE and TP, suggesting the importance especially of openness to trade in matching the observed variation in relative wages (though again the reverse-regression slopes are higher for TP and AC than for P-RE). In all three time windows, moreover, the variance ratios are higher with the P-RE specification than with the TP and AC specifications. The slopes are steeper for the 10-year than for the 5-year window, while for the 14-year window, the slopes are insignificantly different from zero in the full sample, but positive in the restricted sample under the P-RE and TP specifications. These findings are again robust to the choice of the reference sector, as shown in Figures A.6-A.8 in the Online Appendix.[32],[33]

---

[32]We also performed the slope and variance-ratio tests on manufacturing and services separately, treating their labour forces as specific to (and immobile between) them. Cross-section and panel results available upon request are qualitatively similar to the ones obtained on the full sample. They also do not reveal any systematic and robust difference across the two broad sectors. The EHO wage-skill supply relationship seems to fit better the manufacturing sector in the cross-section specifications, but this difference is not confirmed in the within-country specifications.

[33]As in the cross-country analysis, the slope of the regression of predicted values on actual values decreases with the chosen value of σ (1.67 in the results reported). With the P-RE specification, raising σ from 1 to 3 reduces the slope from 0.48 to 0.19 with a 5-year window, from 0.79 to 0.30 with a 10-year window, and from 0.50 to 0.19 with a 14-year window (dropping outliers). The R^2's do not change with σ.

Table 3: Slope and variance-ratio tests – Panel estimates

Bounds:	P-RE		TP		AC	
	$\tilde{w}$ on w	w on $\tilde{w}$	$\tilde{w}$ on w	w on $\tilde{w}$	$\tilde{w}$ on w	w on $\tilde{w}$
5-year changes						
	0.314***	0.291**	0.286***	0.316**	0.224***	0.392**
	(0.0779)	(0.108)	(0.0698)	(0.119)	(0.0531)	(0.157)
Obs	400	400	400	400	400	400
Within R^2	0.221	0.221	0.215	0.221	0.195	0.218
Variance ratio: $\frac{Var(\tilde{w})}{Var(w)}$						
	0.855		0.712		0.455	
10-year changes						
	0.514***	0.482***	0.480***	0.525***	0.380***	0.641***
	(0.121)	(0.138)	(0.111)	(0.152)	(0.0819)	(0.202)
Obs	200	200	200	200	200	200
Within R^2	0.338	0.310	0.335	0.314	0.314	0.306
Variance ratio: $\frac{Var(\tilde{w})}{Var(w)}$						
	0.662		0.556		0.378	
14-year changes						
	0.0690	0.0758	0.0384	0.0496	-0.0256	-0.0445
	(0.189)	(0.207)	(0.176)	(0.227)	(0.152)	(0.262)
Obs	40	40	40	40	40	40
R^2	0.005	0.005	0.002	0.002	0.001	0.001
14-year changes, without outliers						
	0.319*	0.433**	0.280*	0.467**	0.185	0.416
	(0.163)	(0.201)	(0.145)	(0.227)	(0.124)	(0.265)
Obs	37	37	37	37	37	37
R^2	0.138	0.138	0.130	0.130	0.077	0.077
Variance ratio: $\frac{Var(\tilde{w})}{Var(w)}$						
	0.736		0.599		0.444	

Notes: All 5- and 10-year regressions include country fixed effects and year dummies. Standard errors clustered at the country-level (for 5- and 10-year regressions) and heteroskedasticity-robust (for 14-year regressions) are in parenthesis. The bottom panel excludes CHN, MEX and POL from the regressions. Significant at: *10%, **5%, ***1% level.

Overall, the results of the structural tests are supportive of the EHO model. The EHO reallocation mechanisms play an important role in explaining the observed variation in skill premia across countries and within countries over time. The explanatory power of the theoretical framework seems to derive more from variation in openness to trade (through the influence of market shares on purchaser-price elasticities) than in price-ratio elasticities. This

asymmetry may reflect the theoretically more clear-cut effect of openness than of price-ratio elasticities (as discussed in section II), or it may reflect our stark assumptions of zero-or-all per-unit OCWs (allowing $\eta \in (0,1)$ might bring the model even closer to reality). The structural tests also confirm that skill supplies affect wages through channels not mediated by openness or per-unit costs and that variation in skill supplies is by no means the only cause of variation in relative wages.

VI   Comparing reduced-form and structural elasticities

Both the reduced-form and the structural estimates suggest that variation in openness to trade and in price-ratio elasticities contributes importantly to explaining how the elasticity of the relative wages of skilled workers with respect to skill supplies varies across countries and periods. The two sets of estimates (the reduced-form ones being based on the panel estimates in column (6) of Table 1) are also strongly correlated (across all 600 country-year observations, the correlation coefficient is 0.35 (standard error=0.04)).[34]

The reduced-form and structural estimates, however, are measuring different things. The reduced-form estimates can potentially encompass all determinants of the wage-skill supply elasticity (the whole of equation (11)), while the structural estimates in section V cover only the effects of openness and the price-ratio elasticity (the denominator of equation (11), ϕ_{HL}, assuming the numerator to be unity, as in equation (10)). What these structural estimates omit is the influence of directed technical change and the income elasticity of demand, both of which make the numerator of (11) less than unity. The reduced-form estimates should thus be smaller (less negative) than the structural ones, and the difference should reflect the size of the effect of these omitted mechanisms.

Figure 4(a) confirms this theoretical intuition by plotting the kernel densities of the two elasticity estimates, with vertical lines denoting the respective means.[35] The structural estimates are both much larger (i.e. more negative) and less dispersed than the reduced-form ones, with the mean of the reduced-form estimates being -0.18 (the coefficient on $\ln(v)$ in column (6), Table 1) and that of the structural estimates being -0.67.[36] The latter value is close to the estimates from time-series data for the U.S. of Katz and Murphy (1992) and

---

[34]The estimated wage-skill supply elasticity for each country c and year t is: $\widehat{\beta}_{w-v} = \widehat{\beta}_1 + \widehat{\beta}_3 \times \ln(o) + \widehat{\beta}_5 \times \widetilde{\delta}$; where the 'hat' denotes the "estimated" coefficients.

[35]The structural elasticity for each country and year is the median of elasticities across all possible reference sectors.

[36]We do not isolate at this stage elasticities that are not statistically different from zero. In the reduced-form estimates, 33% of the sample has insignificant elasticities at the 10% level. If we restrict our sample only to significant elasticities, the average reduced-form elasticity is -0.22.

Blankenau and Cassou (2011) after controlling for linear time trends (which they interpret as reflecting technical change, but could include the effects of income elasticity), and to that of Acemoglu and Autor (2011) with a non-linear time trend.

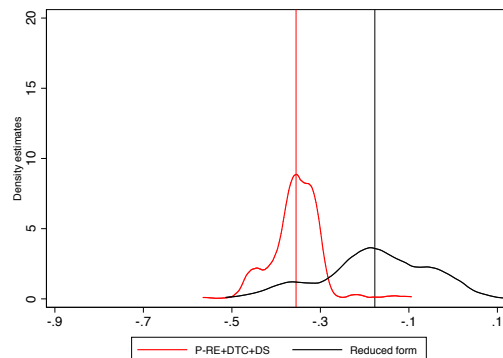Figure 4: Estimated wage-skill supply elasticities

(a) Structural vs. reduced-form
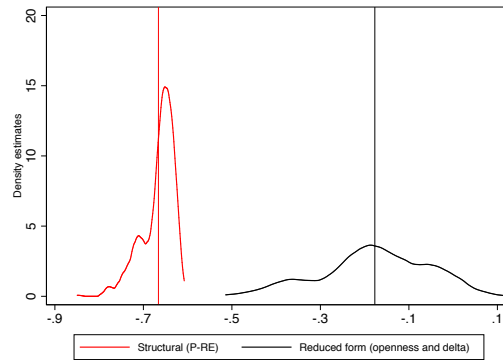
(b) Adding other influences to the structural elasticities
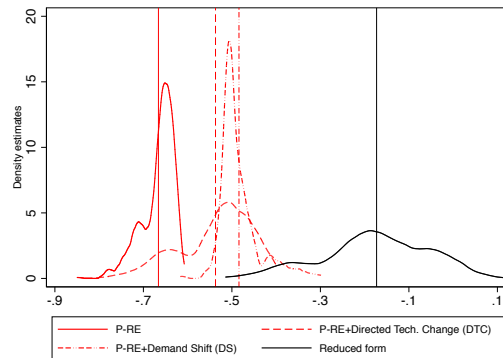
(c) Adding all influences together

Notes: Vertical lines are at the mean elasticities.

The greater dispersion of the reduced-form density is probably due to overestimation of the openness interaction. Equation (10) implies that this coefficient should capture how

greater openness increases the weights of β's relative to γ's in the determination of ϵ's (equation (15)), weakening the inverse relationship between wages and skill supplies. The reduced-form estimate of this coefficient, however, is likely to be biased upward by other influences.

One such influence is the Stolper-Samuelson mechanism in equation (12): the effect of openness on wages varies with a country's skill supplies. Non-HO mechanisms may have a similar effect: for example, Burstein and Vogel (2016, figure 3) show how lower trade costs can cause larger rises in the skilled wage premium in countries with more skilled workers. The openness coefficient also picks up the influence of international trade costs and traded intermediates omitted from our measure of the price-ratio elasticity.

For all these reasons, the wage-skill supply elasticity varies more with openness than in equation (10). Omitting the openness interaction from the calculations (results available upon request) makes the dispersion of the reduced-form density similar to that of the structural density, but has little effect on its mean and hence on the horizontal gap between the distributions.[37]

To account for this gap, the structural estimates need to be extended to include directed technical change (Acemoglu, 2007) and the income elasticity of demand (Caron et al., 2014; Leonardi, 2015), in the ways described by equation (11). Incorporating directed technical change is simple, because it involves only parameters already used for the trade shares and price-ratio elasticity mechanisms.[38] Incorporating the income elasticity of demand requires more parameters to be estimated: π_v, the elasticity of per-capita income with respect to skill supplies; μ_j, the income elasticity of demand for good j; and ω, which captures the effect of income on the skill intensity of the composition of demand for varieties of each good.

Measuring the effect of skill supply on national income has long been seen as difficult, and there is not even a consensus range of values of π_v in the literature (Pritchett, 2001; Hanushek and Woessmann, 2015). For simplicity, we assume $\pi_v = 1$, in line with the emphasis of new growth theory on the externalities associated with human capital. To estimate μ_j, we follow the two-step procedure of Caron et al. (2014), which corrects for differences in prices among countries caused by trade costs (Online Appendix, section A.V). Despite using different data,

---

[37]Another possible explanation for this gap is our benchmark assumption in the structural calculations that $\eta_{j1} = 0$, which is likely to underestimate all the δ_{j1} and thus to overestimate the wage-skill supply elasticity. However, we made the same assumption in the reduced-form estimation, the results of our structural tests were not greatly altered by setting $\delta_{j1} = 1$ (in the TP experiments), and the effects on the mean of including or excluding the price-ratio elasticity term are much the same for the reduced-form and the structural distributions (results available upon request).

[38]Its weakening effect on the wage-skill supply elasticity rests on the assumptions that $\widetilde{\sigma} > 1$ and that $\phi_{HL} > 1$ (which itself increases with $\widetilde{\sigma}$).

we obtain similar results, including the strong correlation of these income elasticities with the skill intensities of the goods concerned. We estimate ω from a reduced-form regression of the skill intensity (skilled/unskilled labour content) of final expenditure on per capita GDP and sector dummies – with both variables in logs and averaged over time within country-sectors.[39] The estimated value of ω is 0.183 (standard error=0.04).

Figures 4(b) and 4(c) plot the resulting kernel densities of the structural wage-skill supply elasticities. Panel (b) adds the directed technical change and income elasticity mechanisms separately: it shows that each of them shifts the mean of the distribution to the right by a similar amount, narrowing the gap between it and the reduced-form distribution, and also that the directed technical change mechanism increases the dispersion of the structural distribution. Panel (c) combines the directed technical change and income elasticity mechanisms. Together, they greatly reduce the gap between the reduced-form mean (-0.18) and the structural mean (now -0.35), as well as creating a substantial overlap between the two distributions.

The gap can be completely eliminated, moreover, by raising the assumed value of σ from 1.67 to 1.92, as shown in the top line of Figure 5 (which spans the generally accepted 1 - 3 σ range). A higher value of σ brings the wage-skill supply elasticity closer to zero in two ways: one is simply to make it easier to substitute skilled for unskilled labour within sectors; the other is to amplify the directedness of technical change by reducing the rate at which the returns to a factor diminish with increases in its effective supply. Without directed technical change, as lines P-RE and P-RE + DS in Figure 5 show, no value of σ could eliminate the gap between the reduced-form and structural means. With directed technical change, by contrast, high values of $\sigma > 2.15$ would make the structural mean positive.

---

[39]The skilled/unskilled labour content of final demand (for both domestically supplied and imported goods) is calculated at the sector level by combining yearly input-output tables and factor use tables from WIOD (following Timmer et al. (2013)). In previous drafts, we estimated sector-specific ω's from panel regressions, some of which however were insignificant. Using those instead of our common ω does not alter the main findings.

Figure 5: Wage-skill supply elasticity under different values of σ
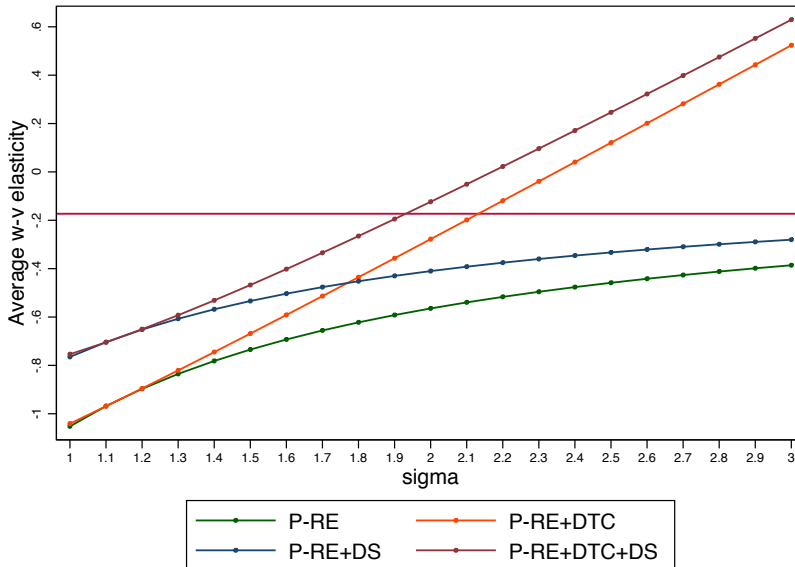
Notes: Horizontal line is at the mean reduced-form elasticity (=-0.177).

These calculations achieve one of the empirical objectives of this paper – to explain the average size of aggregate wage-skill supply elasticities. Our different reduced-form estimates of this average are fairly close together (-0.2 to -0.3 with the cross-country and 14-year change data in Figure 1 and Table 1, and -0.17 in the panel interaction model (column (6) of Table 1)). Our structural estimates show that this average would be three times more negative if it depended only on imperfect substitutability between home and foreign varieties, barriers to trade, and non-proportional costs (the EHO model's openness and price-ratio elasticity mechanisms). Directed technical change and the income elasticity of demand for skill-intensive goods, however, pull the average wage-skill supply elasticity quite strongly towards zero, apparently both to a similar degree (though this conclusion depends, e.g., on our assumption about π_v).

VII Concluding remarks

We have analysed how the relative wages of skilled and unskilled workers in economies open to trade vary with their relative supplies. Our empirical results have shown that, as intuition and much other evidence suggests, wages in open economies vary systematically with skill supplies across countries and over time, but also that, as trade theory suggests, wages are less – sometimes much less – sensitive to skill supplies where barriers to trade are lower.

We have also been able to explain why the average (across countries and periods) wage-skill supply elasticity is about -0.2. That it is negative, rather than zero (as in simple trade theory), is the result of trade barriers, non-proportional costs and imperfect substitutability between home and foreign goods, which make changes in wages necessary to absorb changes in skill supplies. That its negative value is not much greater (in absolute value) and hence closer to standard estimates of the aggregate elasticity of substitution between skilled and unskilled workers reflects the influence of directed technical change and income elasticity in consumption, both of which cause increases in the supply of skilled workers also to increase the demand for skilled workers.

These results are of practical importance, because they make clearer both the scope for and the limitations of policy initiatives to reduce inequality between skilled and unskilled workers by expansion of education and training. As ever, however, there is scope for future improvement of our analysis, especially by using better measures of trade barriers and of per-unit trade and production costs, and by fuller investigation of the effects of trade in intermediate inputs and of payments to capital.

References

Abrego, L., Whalley, J., 2000. Demand Side Considerations and the Trade and Wages Debate. NBER Working Papers 7674.

Acemoglu, D., 2007. Equilibrium Bias of Technology. Econometrica 75(5), 1371–1409.

Acemoglu, D., Autor, D., 2011. Skills, tasks and technologies: Implications for employment and earnings, Elsevier. volume 4 of *Handbook of Labor Economics*, pp. 1043 – 1171.

Acemoglu, D., Gancia, G., Zilibotti, F., 2015. Offshoring and Directed Technical Change. American Economic Journal: Macroeconomics 7(3), 84–122.

Alchian, A., Allen, W., 1964. University Economics. Wadsworth Publishing Company.

Anderson, E., Tang, P.J.G., Wood, A., 2006. Globalization, co-operation costs, and wage inequalities. Oxford Economic Papers 58(4), 569–595.

Anderson, J.E., 1979. A Theoretical Foundation for the Gravity Equation. American Economic Review 69(1), 106–16.

Anderson, J.E., 2011. The Gravity Model. Annual Review of Economics 3(1), 133–160.

Anderson, J.E., Yotov, Y.V., 2010. Specialization: Pro- and Anti-globalizing, 1990-2002. NBER Working Papers 16301.

Arkolakis, C., Costinot, A., Rodriguez-Clare, A., 2012. New Trade Models, Same Old Gains? American Economic Review 102(1), 94–130.

Armington, P.S., 1969. A Theory of Demand for Products Distinguished by Place of Production. Staff Papers - International Monetary Fund 16(1), 159–178.

Autor, D.H., Dorn, D., Hanson, G.H., 2016. The China Shock: Learning from Labor Market Adjustment to Large Changes in Trade. Annual Review of Economics 8(1).

Baldwin, R., Robert-Nicoud, F., 2014. Trade-in-goods and trade-in-tasks: An integrating framework. Journal of International Economics 92(1), 51–62.

Barro, R.J., Lee, J.W., 2013. A new data set of educational attainment in the world, 1950-2010. Journal of Development Economics 104, 184 – 198.

Berman, N., Martin, P., Mayer, T., 2012. How do Different Exporters React to Exchange Rate Changes? The Quarterly Journal of Economics 127(1), 437–492.

Bernard, A.B., Redding, S.J., Schott, P.K., 2013. Testing for Factor Price Equality with Unobserved Differences in Factor Quality or Productivity. American Economic Journal: Microeconomics 5(2), 135–63.

Bernard, A.B., Redding, S.J., Schott, P.K., Simpson, H., 2008. Relative Wage Variation and Industry Location in the United Kingdom. Oxford Bulletin of Economics and Statistics 70(4), 431–459.

Blankenau, W., Cassou, S., 2011. Industry estimates of the elasticity of substitution and the rate of biased technological change between skilled and unskilled labour. Applied Economics 43(23), 3129–3142.

Blum, B.S., 2010. Endowments, Output, and the Bias of Directed Innovation. The Review of Economic Studies 77(2), 534–559.

Broda, C., Weinstein, D.E., 2006. Globalization and the Gains from Variety. The Quarterly Journal of Economics 121(2), 541–585.

Burstein, A., Cravino, J., Vogel, J., 2013. Importing skill-biased technology. American Economic Journal: Macroeconomics 5(2), 32–71.

Burstein, A., Hanson, G., Tian, L., Vogel, J., 2017. Tradability and the Labor-Market Impact of Immigration: Theory and Evidence from the U.S. Technical Report 23330. National Bureau of Economic Research.

Burstein, A., Vogel, J., 2016. International Trade, Technology and the Skill Premium. Journal of Political Economy Forthcoming.

Caron, J., Fally, T., Markusen, J.R., 2014. International Trade Puzzles: A Solution Linking Production and Preferences. The Quarterly Journal of Economics 129(3), 1501–1552.

Caselli, F., Feyrer, J.D., 2007. The Marginal Product of Capital. The Quarterly Journal of Economics 122(2), 535–568.

Chor, D., 2010. Unpacking Sources of Comparative Advantage: A Quantitative Approach. Journal of International Economics 82(2), 152–167.

Ciccone, A., Peri, G., 2005. Long-Run Substitutability Between More and Less Educated Workers: Evidence from U.S. States, 1950-1990. The Review of Economics and Statistics 87(4), 652–663.

Costinot, A., Rodriguez-Clare, A., 2014. Trade Theory with Numbers: Quantifying the Consequences of Globalization, in: Elhanan Helpman, K.R., Gopinath, G. (Eds.), Handbook of International Economics. Elsevier. volume 4, pp. 197 – 261.

Davis, D.R., Weinstein, D.E., 2001. An Account of Global Factor Trade. American Economic Review 91(5), 1423–1453.

Dornbusch, R., 1987. Exchange rates and prices. American Economic Review 77(1), 93–106.

Eaton, J., Kortum, S., 2002. Technology, Geography, and Trade. Econometrica 70(5), 1741–1779.

Epifani, P., Gancia, G., 2006. Increasing Returns, Imperfect Competition, and Factor Prices. The Review of Economics and Statistics 88(4), 583–598.

Epifani, P., Gancia, G., 2008. The skill bias of world trade. Economic Journal 118(530), 927–960.

Feenstra, R.C., 1994. New Product Varieties and the Measurement of International Prices. American Economic Review 84(1), 157–77.

Feenstra, R.C., Luck, P., Obstfeld, M., Russ, K.N., 2017. In Search of the Armington Elasticity. The Review of Economics and Statistics Forthcoming.

Freeman, R.B., 1995. Are Your Wages Set in Beijing? Journal of Economic Perspectives 9(3), 15–32.

Grossman, G.M., Rossi-Hansberg, E., 2008. Trading Tasks: A Simple Theory of Offshoring. American Economic Review 98(5), 1978–97.

Hanson, G.H., Slaughter, M.J., 2002. Labor-market Adjustment in Open Economies: Evidence from US states. Journal of International Economics 57(1), 3–29.

Hanushek, E.A., Woessmann, L., 2015. The Knowledge Capital of Nations: Education and the Economics of Growth. MIT Press.

Harrison, A., McLaren, J., McMillan, M., 2011. Recent Perspectives on Trade and Inequality. Annual Review of Economics 3(1), 261–289.

Haskel, J., Lawrence, R.Z., Leamer, E.E., Slaughter, M.J., 2012. Globalization and U.S. Wages: Modifying Classic Theory to Explain Recent Facts. Journal of Economic Perspectives 26(2), 119–40.

Helpman, E., Itskhoki, O., 2010. Labour Market Rigidities, Trade and Unemployment. Review of Economic Studies 77(3), 1100–1137.

Hijzen, A., Swaim, P., 2010. Offshoring, labour market institutions and the elasticity of labour demand. European Economic Review 54(8), 1016–1034.

Hummels, D., Skiba, A., 2004. Shipping the Good Apples Out? An Empirical Confirmation of the Alchian-Allen Conjecture. Journal of Political Economy 112(6), 1384–1402.

Johnson, G., Stafford, F., 1999. The labor market implications of international trade, Elsevier. volume 3 of *Handbook of Labor Economics*, pp. 2215 – 2288.

Jones, R.W., 1965. The Structure of Simple General Equilibrium Models. Journal of Political Economy 73(6), 557–572.

Katz, L.F., Autor, D.H., 1999. Changes in the Wage Structure and Earnings Inequality, volume 3 of *Handbook of Labor Economics*. chapter 26, pp. 1463–1555.

Katz, L.F., Murphy, K.M., 1992. Changes in Relative Wages, 1963-1987: Supply and Demand Factors. The Quarterly Journal of Economics 107(1), 35–78.

Kiyota, K., 2011. Paths of Development and Wage Variations. Review of International Economics 19(4), 697–717.

Kiyota, K., 2012. A Many-cone World? Journal of International Economics 86(2), 345–354.

Klepper, S., Leamer, E.E., 1984. Consistent Sets of Estimates for Regressions with Errors in All Variables. Econometrica 52(1), 163–83.

Krugman, P.R., 1979. Increasing returns, monopolistic competition, and international trade. Journal of International Economics 9(4), 469–479.

Lafortune, J., Tessada, J., González-Velosa, C., 2015. More hands, more power? estimating the impact of immigration on output and technology choices using early 20th century us agriculture. Journal of International Economics 97(2), 339 – 358.

Leamer, E.E., 2012. The Craft of Economics: Lessons from the Heckscher-Ohlin Framework. The Ohlin Lectures, MIT Press.

Leonardi, M., 2015. The Effect of Product Demand on Inequality: Evidence from the United States and the United Kingdom. American Economic Journal: Applied Economics 7(3), 221–47.

Lewis, E., 2013. Immigration and Production Technology. Annual Review of Economics 5(1), 165–191.

Markusen, J.R., Venables, A.J., 2007. Interacting factor endowments and trade costs: A multi-country, multi-good approach to trade theory. Journal of International Economics 73(2), 333–354.

Marshall, K.G., 2012. International productivity and factor price comparisons. Journal of International Economics 87(2), 386 – 390.

Melitz, M.J., 2003. The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity. Econometrica 71(6), 1695–1725.

Morrow, P., Trefler, D., 2017. Endowments, Factor Prices, and Skill-Biased Technology: Importing Development Accounting into HOV. Technical Report January.

OECD, 2013. OECD Indicators of Employment Protection. Technical Report. Organisation for Economic Cooperation and Development (OECD).

Panagariya, A., 1999. Trade Openness: Consequences for the Elasticity of Demand for Labor and Wage Outcomes. International Trade 0308007. EconWPA.

Patel, N., Wang, Z., Wei, S.J., 2014. Global Value Chains and Effective Exchange Rates at the Country-Sector Level. NBER Working Papers 20236.

Pritchett, L., 2001. Where Has All the Education Gone? The World Bank Economic Review 15(3), 367–391.

Ramondo, N., Rodrguez-Clare, A., Saboro-Rodrguez, M., 2016. Trade, Domestic Frictions, and Scale Effects. American Economic Review 106(10), 3159–84.

Robbins, D.J., 1996. Evidence on Trade and Wages in the Developing World. OECD Development Centre Working Papers 119. OECD Publishing.

Rodrik, D., 1997. Has Globalization Gone Too Far? Institute for International Economics, Washington D.C.

Romalis, J., 2004. Factor Proportions and the Structure of Commodity Trade. American Economic Review 94(1), 67–97.

Sato, K., 1967. A two-level constant elasticity-of-substitution production function. The Review of Economic Studies 34(2), 201–218.

Schott, P.K., 2003. One Size Fits All? Heckscher-Ohlin Specialization in Global Production. American Economic Review 93(3), 686–708.

Senses, M., 2010. The effects of offshoring on the elasticity of labor demand. Journal of International Economics 81(1), 89–98.

Slaughter, M.J., 2001. International trade and labor-demand elasticities. Journal of International Economics 54(1), 27–56.

Soderbery, A., 2015. Estimating import supply and demand elasticities: Analysis and implications. Journal of International Economics 96(1), 1 – 17.

Timmer, M.P., 2012. The World Input-Output Database (WIOD): Contents, Sources and Methods. Technical Report. WIOD Working Paper Number 10.

Timmer, M.P., Erumban, A.A., Los, B., Stehrer, R., de Vries, G.J., 2014. Slicing up global value chains. Journal of Economic Perspectives 28(2), 99–118.

Timmer, M.P., Los, B., Stehrer, R., Vries, G.J., 2013. Fragmentation, incomes and jobs: an analysis of European competitiveness. Economic Policy 28(76), 613–661.

Tomiura, E., 2005. Factor Price Equalization In Japanese Regions. The Japanese Economic Review 56(4), 441–456.

Trefler, D., 1995. The Case of the Missing Trade and Other Mysteries. The American Economic Review 85(5), 1029–1046.

Trionfetti, F., 2015. Comparative skill premia. mimeo.

Wood, A., 1994. North-South Trade, Employment and Inequality: Changing Fortunes in a Skill-Driven World. Oxford University Press.

Wood, A., 2002. Globalization and wage inequalities: A synthesis of three theories. Review of World Economics (Weltwirtschaftliches Archiv) 138(1), 54–82.

Wood, A., 2012. Better-behaved Heckscher-Ohlin models through more consistent treatment of trade costs. QEH Working Paper Series, n. 185. URL: http://www3.qeh.ox.ac.uk/RePEc/qeh/qehwps/qehwps185.pdf.