



Unobservable skill dispersion and comparative advantage[☆]



Matilde Bombardini^{a,b,c,*}, Giovanni Gallipoli^a, Germán Pupato^d

^a University of British Columbia, Canada

^b CIFAR, Canada

^c NBER, United States

^d Ryerson University, Canada

ARTICLE INFO

Article history:

Received 31 March 2012

Received in revised form 8 November 2013

Accepted 12 November 2013

Available online 23 November 2013

JEL classification:

F12

F16

J82

Keywords:

Substitutability

Skill dispersion

Skill distribution

ABSTRACT

This paper investigates a theoretical mechanism linking comparative advantage to the distribution of skills in the working population. We develop a tractable multi-country, multi-industry model of trade with unobservable skills in the labour market and show that comparative advantage derives from (i) cross-industry differences in the substitutability of workers' skills and (ii) cross-country differences in the dispersion of skills. We establish the conditions under which higher skill dispersion leads to specialization in industries characterized by higher skill substitutability across tasks. The main results are robust when the model is extended to allow for partial observability of skills. Finally, we use distributions of literacy scores from the International Adult Literacy Survey to approximate cross-country productivity differences due to skill dispersion and we carry out a quantitative assessment of the impact of skill dispersion on the pattern of trade.

© 2013 Elsevier B.V. All rights reserved.

1. Introduction

The theory of comparative advantage identifies factor endowments as a key determinant of the pattern of trade. In particular, the theoretical prediction that countries endowed with larger stocks of human capital export relatively more in skill-intensive industries has received support in the literature, see Romalis (2004). In previous work (Bombardini et al., 2012b, henceforth BGP) we build on this line of research and argue that the second moment of the distribution of skills also determines comparative advantage. In particular, we find that the degree of skill dispersion has a quantitative impact on trade flows similar to that of the aggregate endowment of human capital. This paper presents a multi-country multi-industry model that shows how skill dispersion generates comparative advantage and thus provides a theoretical underpinning to the empirical evidence in BGP.

Why would the skill distribution matter for specialization and trade? We argue that industries vary in the degree of substitutability of workers' skills in the production process. In particular, some industries, such as aerospace or engine manufacturing, require completing long sequences of tasks and poor performance at any single stage greatly reduces the value of output. These are industries with low skill substitutability, or O-ring as in Kremer (1993), where efficiency improves when workers of similar skills are employed in every stage of production. In other industries, such as paper manufacturing, teamwork is relatively less important, as skills are more easily substitutable and poor performance in some task can be mitigated by superior performance in others.

We investigate theoretically whether countries with greater skill dispersion specialize in industries characterized by higher substitutability of skills across tasks. We build a model with many countries and many industries. Countries only differ in the distribution of skills in the labour force, while industries only differ in the degree of skill substitutability in the production process. At the micro level, our framework features worker heterogeneity and non-linear returns to scale; however, at the industry level, it is isomorphic to a perfectly competitive model with CRS and technological differences across many countries and industries, as in Costinot et al. (2012).

We develop the central argument in a benchmark framework in which skills are not observable by firms, adapted from Akerlof (1970). Skill unobservability is not an unusual assumption in the literature,

[☆] We would like to thank Paul Beaudry, Pablo D'Erasmus, David Green, Patrick Francois, Keith Head, Thomas Lemieux, Giovanni Maggi, Vadim Marmer, Thomas Sampson, Francesco Trebbi, Jonathan Vogel and seminar participants at CIFAR, EIT, Getulio Vargas Foundation, Queen's University, Ryerson University, University of Alberta, University of Auckland, UC Davis, UC San Diego, University of Chicago Booth, Universidad de Chile, University of Sydney and Yale University for the helpful comments. Bingjing Li and Ruoying Wang provided excellent research assistance.

* Corresponding author at: 997-1873 East Mall, Vancouver, BC V6T 1Z1, Canada.

E-mail address: matilde.bombardini@ubc.ca (M. Bombardini).

see for example Grossman and Maggi (2000), Grossman (2004) and Helpman et al. (2010).¹ This modelling choice is motivated by the facts, documented in BGP, that (i) observable characteristics of workers, including age and education, account for a minor share of total variation in work-related literacy scores within countries; and (ii) groups of observationally similar workers exhibit very different degrees of skill dispersion across countries.² One immediate advantage of this approach is tractability in a setting with many countries and many sectors. Essentially, under skill unobservability, every worker is indifferent between employment in any two firms because there is a single (type-independent) wage that clears the labour market. In equilibrium, workers are randomly matched to firms and thus the latter inherit the skill distribution prevailing in the economy – i.e. the distribution of workers' unobservable skills in every firm is identical to the distribution of unobservable skills in the country.

The results in the benchmark model carry over with minor qualifications when cross-country differences in observable skills are introduced in the analysis. When firms can observe some components of skills, the equilibrium features sorting on observable skills by industry. As in the benchmark model, however, we show that unobservable skill dispersion fully determines the pattern of trade in equilibria with wage equalization across countries. In light of this result, the benchmark model is best interpreted as a mechanism illustrating how the dispersion of skills among workers with otherwise identical observable characteristics affects comparative advantage. In the rest of the paper, we thus sometimes refer to unobservable skills as 'residual' skills.

We show that the interaction of skill dispersion and skill substitutability generates differences in output per worker across industries and countries, driving the pattern of international trade. The central result of the paper establishes conditions under which countries with a high dispersion of skills in the labour force are relatively more productive, and thus export relatively more, in sectors where skills are more easily substitutable across tasks. Interestingly, we also show that the effects of skill dispersion on output and specialization are isomorphic to the effects of technological differences in Ricardian models, such as Costinot et al. (2012). In this sense, our work has implications for the quantitative assessment of Ricardian comparative advantage since cross-country differences in measured total factor productivity can arise as the by-product of differences in the distribution of factor endowments.³

The theoretical framework developed in this paper can also be used for quantitative analysis. Using distributions of literacy scores from the International Adult Literacy Survey (IALS), we construct productivity differences attributable to skill dispersion. Cross-country differences in trade costs are calibrated to match bilateral trade flows at the industry level. We use the numerical counterpart of the model to carry out a quantitative assessment of the impact of skill dispersion on trade patterns. In particular, we measure the general equilibrium responses in industry-level trade when cross-country differences in skill dispersion are eliminated. These experiments, which we run under a variety of technology parametrizations, suggest a significant role of skill dispersion on trade flows, with different effects on different countries.

¹ In Grossman and Maggi (2000) and Helpman et al. (2010) firms can observe a component of individual skills before hiring workers. There remains, however, another component which firms never observe. We introduce a similar extension in Section 6. In Grossman (2004), as in this paper, individual skills are unobservable by firms and the production process features imperfect skill substitutability.

² Naturally, firms are likely to be able to observe the skills of their employees more accurately than econometricians. Therefore, these facts should be interpreted as playing a motivational role for our modelling approach, and not as precise descriptions of the degree of uncertainty that firms face regarding the skills of their employees.

³ Costinot et al. (2012) argue that their exogenous productivity differences aim to capture factors such as climate, infrastructure, and institutions, which affect the productivity of all producers in a given country and industry. To the extent that the distribution of human capital endowments – above and beyond observable credentials – is the product of a country's social structure and norms, our explanation would pertain to the institutional view of comparative advantage.

This paper is related to recent theoretical research studying how skill distributions influence the pattern of trade. The hypothesis that skill dispersion may lead to specialization was first put forth by Grossman and Maggi (2000) in the context of a two-country, two-sector model, with competitive labour markets and constant returns to skills. They show that, when skills are fully observable, cross-country differences in skill dispersion do not generate comparative advantage when production technologies display convex isoquants in worker skills.⁴ Gains from trade do not exist because, in equilibrium, workers of identical abilities are paired together, i.e. self-matching prevails in every industry, making skill dispersion irrelevant.⁵ Grossman and Maggi (2000) also consider the case in which a portion of individual skills is unobservable. As in our paper, workers (with identical observable skills) are randomly matched to firms. Although they show that skill dispersion generates gains from trade when both industries have convex isoquants, they cannot determine the resulting *pattern* of trade across countries, which is the main focus of this paper.⁶ In order to make progress, we propose a specification of production technologies that gives us sufficient tractability to handle a wide range of substitution elasticities, i.e. varying degrees of concavity or convexity of the isoquants, in a parametric way. In combination with random matching, this assumption allows us to characterize the variation in productivity differentials arising from skill dispersion, across multiple countries and industries, which is the crucial step in pinning down the pattern of trade.

Interest in the relevance of skill distributions for trade is relatively recent. Ohnsorge and Trefler (2007) propose a Roy-type model with two-dimensional worker heterogeneity to show that, when each worker represents a bundle of two skills, the correlation of the two in the population determines comparative advantage. Grossman (2004) starts from the premise that, in some sectors, incomplete contracts make it difficult to tie remuneration to an individual worker's output. In a country with high skill dispersion highly skilled individuals prefer to sort into sectors where individual performance is easier to measure, rather than working in an industry where the common wage is dragged down by workers with relatively low skills. This type of endogenous sorting results in comparative advantage. Finally, in Bougheas and Riezman (2007) comparative advantage emerges from differential returns to skills across sectors.

The next section describes consumer preferences, production technologies and the labour market. Section 3 discusses how unobservable skill dispersion generates productivity differences across countries and industries, driving comparative advantage. Section 4 studies the optimization problem of individual firms. Section 5 analyzes the implications of skill dispersion for the pattern of international trade. Section 6 introduces observable skills in the model. Section 7 presents the quantitative analysis and our counterfactual results. The paper ends with some concluding remarks.

2. Setup

2.1. Preferences

Countries are denoted by a subscript $c \in \{1, \dots, C\}$, which is dropped when it creates no ambiguity. Each country c is populated by a measure L_c of individuals. Utility of the representative consumer depends on the

⁴ Supermodularity implies that the marginal product of any worker is increasing in the ability of the co-worker. Submodularity of the production function implies the opposite.

⁵ In the case of observable skill dispersion, gains from trade are conditional on the existence of a supermodular sector, where self-matching prevails, and of a submodular sector, where the most skilled workers are paired with the least skilled co-workers, i.e. cross-matching prevails. In this case, the country with a more dispersed skill distribution specializes in the submodular sector.

⁶ When one sector displays convex isoquants and the other displays concave isoquants, they show that unobservable skill dispersion reinforces the pattern of trade induced by observable skill dispersion.

consumption of differentiated goods Q_i with $i \in \{1, \dots, I\}$. The utility function U is Cobb–Douglas:

$$\log U = \sum_{i=1}^I \alpha_i \log Q_i,$$

where $0 < \alpha_i < 1$, $\sum \alpha_i = 1$ and Q_i is an aggregate consumption index over a fixed set Ω_i of varieties of i . Preferences exhibit constant elasticity of substitution σ_i across varieties of any good.⁷ As a result, total expenditure on variety ω of good i is:

$$x_i(\omega) = \left[\frac{p_i(\omega)}{P_i} \right]^{1-\sigma_i} \alpha_i E, \tag{1}$$

where E is the aggregate expenditure, $p_i(\omega)$ is the price of variety ω of i , and P_i is the ideal CES price index of Q_i .

2.2. Production

Each variety ω in the differentiated industry i is produced under perfect competition and free entry. The typical firm producing ω has to incur a fixed start-up cost of f units of the numeraire good, to be specified below. The amount of output produced $y(\omega)$ depends on the skill level of each worker hired $a < 0$, the measure of workers hired h , the distribution of skills across workers $\tilde{G}(a)$ and a random productivity shock $\varepsilon > 0$. The distribution of skills matters for production because we assume that skills are imperfectly substitutable. In particular, the production function depends on the degree of substitutability λ among workers' skills in that industry and takes the form:

$$y = \varepsilon \left(\int h a^\lambda d\tilde{G}(a) \right)^\frac{\gamma}{\lambda} \text{ with } 0 < \gamma < \lambda. \tag{2}$$

The parameter $\lambda \in \Lambda$ measures the degree of skill substitutability. The elasticity of substitution among skill levels, for a fixed mass of workers h , is given by $\frac{1}{1-\lambda}$. The set $\Lambda \subset (0, 1]$ determines the admissible range of substitutability. A key assumption is that each industry i is characterized by a different value of λ in production, and therefore by a different degree of substitutability among workers' skills. Since λ is the only characteristic that distinguishes technology across industries,⁸ in the remainder of the paper we drop subscript i and index industries by their corresponding parameter λ .

Note that when skills are perfectly substitutable (that is, $\lambda = 1$) output depends on the average skills of workers but not on higher moments of the skill distribution and skill dispersion does not generate comparative advantage. In what follows, we exploit this natural property of the model to enhance the tractability of the analysis. In particular, we show that, if there is an industry with no transport costs, $\lambda = 1$ and no intra-industry goods' heterogeneity, equilibria with incomplete specialization display wage equalization across countries.⁹ Throughout we assume that $1 \in \Lambda$ and use Λ_{-1} to denote the set of industries with imperfect skill substitutability; i.e. $\Lambda_{-1} = \Lambda \setminus \{1\}$.

In order to obtain a clear-cut characterization of the impact of skill dispersion on bilateral trade flows at the industry level, we introduce intra-industry heterogeneity in industries with imperfect skill substitutability, as in Costinot et al. (2012). In particular, for industries $\lambda \in \Lambda_{-1}$, the Ricardian productivity shock $\varepsilon = \varepsilon(\omega)$ is assumed to be an i.i.d. variety-specific draw from a Fréchet distribution $F(\varepsilon) = \exp[-\varepsilon^{-\theta}]$

for $\varepsilon > 0$, with $\theta > 1$. As mentioned, we assume that the distribution of ε is degenerate when skills are perfectly substitutable and set $\varepsilon(\omega) = 1$ for all varieties in industry $\lambda = 1$. Importantly, these Ricardian shocks are identical across countries and therefore do not generate comparative advantage in the model.

Two properties of this production function are worth mentioning. First, every worker has a positive marginal product: in particular, the marginal product of adding worker h of skill level a_h is $\varepsilon h^{\frac{\gamma}{\lambda}-1} \left(\int a^\lambda d\tilde{G}(a) \right)^{\frac{\gamma}{\lambda}-1} \frac{\gamma}{\lambda} a_h^\lambda$. Second, the parameter γ controls the return to the mass of workers, given λ and the skill distribution. In this setting with price taking and fixed costs, decreasing marginal returns in every industry are necessary for the existence of an equilibrium with free entry.¹⁰ In what follows, we restrict Λ to ensure $\gamma < \lambda$.¹¹

2.3. Labour market

We model an environment with informational asymmetries, adapted from Akerlof (1970). Workers are characterized by different skill levels. Skill is a continuous variable distributed in the population of country c according to a distribution function $G_c(a)$ with support $[a_{\min}, a_{\max}]$, where $a_{\max} > a_{\min} > 0$. Importantly, skills are unobservable by the firms. A competitive equilibrium thus consists of a single wage rate w_c , independent of worker types, that clears the labour market.

Since the opportunity cost of accepting employment is the same for every worker, decisions to participate in the labour market are independent of the wage rate and a worker's type. Moreover, since every firm offers the same wage, it follows that every worker is indifferent between any two jobs in the economy. As a result, there is no sorting between workers and firms in equilibrium. This implies that every firm, in any industry λ and country c , inherits the residual skill distribution in the labour force:

$$\tilde{G}(a) = G_c(a).$$

There are a few considerations to make about random matching. First, as will become clear below, random matching allows us to study the essential implications of unobservable skills for comparative advantage within a tractable general equilibrium framework. More sophisticated forms of imperfect matching would result if screening or signalling technologies were introduced into the model. In general, these extensions would come at the cost of analytic solutions. However, our analysis in Section 6 can be interpreted as an attempt to partially relax these assumptions while keeping sufficient tractability. When worker skills are modelled as the product of two components, one of which is observable by firms, we can think of either (i) firms having access to a costless screening technology that allows them to identify some specific dimension of worker skills or, alternatively, (ii) workers that can credibly and costlessly signal some of their skills. Second, our focus is mainly on unobservable skills because this case has received relatively less attention in the literature and, as we show below, it can generate patterns of comparative advantage that are consistent with

⁷ More specifically: $Q_i = \left[\int_{\omega \in \Omega_i} q_i(\omega)^{\frac{\sigma_i-1}{\sigma_i}} d\omega \right]^{\frac{\sigma_i}{\sigma_i-1}}$ with $\sigma_i > 1$ where $q_i(\omega)$ is the quantity consumed of variety ω of good i .

⁸ The parameter γ is constant across industries and does not affect comparative advantage.

⁹ The industry $\lambda = 1$ therefore plays a similar role to a homogeneous sector in many trade models.

¹⁰ Note that the return to the mass of workers also depends on λ and therefore varies across industries. However, as the analysis in Section 5 shows, this effect plays no role in shaping comparative advantage in equilibria with wage equalization across countries.

¹¹ Relative to Grossman and Maggi (2000), here the emphasis is on the degree of substitutability across sectors and not on whether the production function is submodular or supermodular. In fact, because we introduce sufficiently strong decreasing returns, our production function is always submodular. We should stress though that this is inconsequential for the case of unobservable skills. The crucial factor is the degree of substitutability across sectors, which, in the case of only two tasks would be easily represented graphically by the curvature of the isoquants. Notice that, differently from Grossman and Maggi (2000), here isoquants may be convex even with a submodular production function (this is possible when we remove the crucial assumption of constant returns to talent present in Grossman and Maggi (2000)). This topic is explored in Bombardini et al. (2012a).

the empirical evidence in BGP. In Section 6, we show that the results of this benchmark setup are robust when cross-country differences in observable skills are introduced in the analysis.

3. Skill dispersion as comparative advantage

This section shows how cross-country differences in the distribution of residual skills generate comparative advantage. To facilitate the discussion we write the production function in Eq. (2) as $y = \varepsilon A(\lambda, c) h^{\frac{\lambda}{1-\lambda}}$ where the factor $A(\lambda, c)$ is defined as:

$$A(\lambda, c) \equiv \left(\int a^\lambda dG_c(a) \right)^{\frac{1}{\lambda}} \tag{3}$$

We refer to $A(\lambda, c)$ as ‘fundamental productivity’, although it is not the result of countries having access to different technologies. The magnitude of $A(\lambda, c)$ depends on the combination of a country-specific skill distribution and an industry-specific level of skill substitutability in production. Therefore, unlike the i.i.d. technological shocks captured by ε , fundamental productivity varies systematically across countries and industries and, as Section 5 shows, it is the sole determinant of the pattern of comparative advantage in this model. The goal of this section is to understand how variation in $A(\lambda, c)$ is affected by the distribution of skills.¹²

Motivated by the empirical evidence presented in BGP, we explore the conditions under which countries with higher skill dispersion have a comparative advantage in sectors with higher skill substitutability. Property 1, stated below, provides a general condition for this pattern of comparative advantage to emerge from differences in the distribution of skills.¹³

For simplicity we compare countries that have the same average skills, but different dispersion. We order countries so that, if $c < c'$, then country c' is characterized by a skill distribution $G_{c'}(a)$ which is a mean-preserving spread of the skill distribution $G_c(a)$ in country c .

Property 1. $A(\lambda, c)$ is strictly log-supermodular in λ and c , i.e. for any $\lambda, \lambda' \in \Lambda$ and $c, c' \in \{1, \dots, C\}$ such that $\lambda < \lambda'$ and $c < c'$:

$$\frac{A(\lambda, c')}{A(\lambda, c)} < \frac{A(\lambda', c')}{A(\lambda', c)} \tag{4}$$

Property 1 states that the fundamental productivity of firms in countries with high skill dispersion will be relatively larger in high substitutability sectors.

A general result of this type cannot be established for all skill distributions. We therefore provide three different approaches to studying this problem. First, we show that comparative advantage can be established for any distribution with an upper-bounded support by appropriately restricting the admissible range of substitutability. Second, we perform comparative statics for specific distributions of skills that do not satisfy the sufficient conditions of the first approach. Finally, we use the empirical distributions of IALS scores in 19 countries to check numerically whether Property 1 holds when λ is allowed to vary over a wider range, relaxing the condition $\lambda \leq 1$.

Our first approach yields a general result for skill distributions with bounded supports that relies on restricting the admissible range of skill substitutability.¹⁴




Fig. 1. Admissible substitutability for bounded skill distributions.

Proposition 1. Property 1 holds under the following sufficient conditions:

- (i) the skill distribution in any country is supported on a bounded interval $[a_{\min}, a_{\max}]$, where $a_{\max} > a_{\min} > 0$;
- (ii) the set of admissible substitutability satisfies $\Lambda \subset [\lambda_{\min}, 1]$ where $0 < \lambda_{\min} < 1$, implicitly defined by

$$\log a_{\max} = \frac{2\lambda_{\min} - 1}{(1 - \lambda_{\min})\lambda_{\min}}$$

Fig. 1 displays the admissible range of substitutability for an arbitrary skill distribution $G_c(a)$ with bounded support. The lower bound on substitutability is a function of the upper bound of the skill distribution, $\lambda_{\min}(a_{\max})$, as defined in Proposition 1. As a_{\max} increases, the lower bound on the admissible substitutability restriction becomes tighter, meaning that comparative advantage can be established for sets of industries with relatively higher skill substitutability in production. This means that Property 1 will not hold for some highly concave functions if the distribution of skills has a sufficiently large upper bound, a somewhat counterintuitive result which we further discuss at the end of this section.

Our second approach to studying Property 1 relaxes the conditions on substitutability at the cost of concentrating on specific parametric distributions. We consider continuous distributions that are characterized by at least two parameters (in order to be able to consider mean-preserving increases in dispersion) and are defined on a positive support. In the appendix, we show that if skills are distributed according to a Log-normal distribution then, if country c and c' are characterized by skill distributions $G_c(a)$ and $G_{c'}(a)$ such that $G_{c'}(a)$ has equal mean and higher variance than $G_c(a)$ and if $\lambda < \lambda'$ then Property 1 holds. A similar result is established for the Pareto distribution. Besides these analytical results, we have also numerically computed the A 's for several other distributions, including the uniform, triangular, gamma, beta and inverse Gaussian. For all these distributions, and for a wide range of parameters, it was not possible to find a numerical violation of the ranking in Eq. (4).

These analytical results provide sound support for Eq. (4) under the condition $\lambda \leq 1$. However, in anticipation of our quantitative exercises in Section 7, where some sectors have a calibrated λ larger than one, we check numerically whether Property 1 is verified over a wider range of λ . In particular, we extend a numerical exercise originally presented in

¹² Although in a different context, Benabou (2005) analyzes the related issue of endogenous technological choice as a function of the degree of skill dispersion in the population.

¹³ It is important to notice that the patterns of comparative advantage and trade are always fully determined in this model once skill dispersion and substitutability have pinned down an arbitrary $A(\lambda, c)$ function, regardless of whether Property 1 holds or not.

¹⁴ Imposing some upper bound on a is a reasonable restriction, as it only rules out the existence of infinitely productive workers.

BGP, using the empirical distribution of IALS test scores in 19 countries to construct $A(\lambda, c)$ for an equally spaced grid of 50 different λ s defined over the $[0.1, 5]$ interval.¹⁵ The data are described in Section 7. This grid covers the full range of λ s that we later employ in the quantitative section of the paper.¹⁶ We obtain that, when averaging across country pairs, $\frac{A(\lambda, c)}{A(\lambda, c)}$ is increasing in λ at 96.9% of the grid points, where c is a country with higher skill dispersion than c (as measured by the coefficient of variation of IALS scores). The corresponding exercise on the $[0.1, 1]$ interval yields a similar result, with relative productivity increasing at 97.3% of the grid points.

As a robustness check of the relevance of Property 1 we have investigated under what conditions a theoretical violation of the ranking could be engineered. One might expect that increasing differences in substitutability between production technologies would result in more, rather than less, relative advantage. Instead, Proposition 1 implies that Property 1 will not hold if some industries have low enough substitutability in production. To make sense of this result we rely on findings pertaining to choice theory and risk aversion, due to Ross (1981). He shows that, if one adopts the Arrow–Pratt definition of risk aversion (essentially a measure of concavity), there exist lotteries such that a more risk-averse individual may be willing to pay less than a less risk-averse individual in order to avoid an increase in risk. The specific lottery employed as an example by Ross can be explained in the context of our model if one reinterprets $A(\lambda, c)^{\frac{1}{\gamma}}$ as certainty equivalent, i.e. the constant skill level that would make a firm as productive, and $f(a) = a^\lambda$ as utility function. Starting from a skill distribution where most of the workers have low skills, with the exception of few very talented individuals, consider adding a small amount of dispersion at the high skill level. A sector with lower λ may counterintuitively see its certainty equivalent drop by relatively less with such increase in dispersion. The intuition is that the increase in dispersion at very high skills happens in a range where $f(a)$ has relatively little curvature. Moreover, in order to avoid an increase in dispersion, a firm in a low λ sector would have to consider lowering its certainty equivalent in a relatively steep portion of $f(a)$ and could be less willing to do so (relative to a higher λ sector). Using the example by Ross it is possible to find a small set of parameters which generate such a violation of Property 1. Proposition 1 spells out a condition which rules out this phenomenon, by lowering the upper bound of skills to the point at which higher skill dispersion always results in large enough output losses.

4. The firm's problem

This section analyzes the optimal entry and employment choices of a typical firm producing variety ω of good λ in country c . Under perfect competition, the price $p_j(\omega, \lambda)$ that consumers in country j pay for variety ω of good λ is

$$p_j(\omega, \lambda) = \min_{1 \leq c \leq C} \{ \mu_{cj}(\omega, \lambda) \},$$

where $\mu_{cj}(\omega, \lambda)$ is the unit cost of producing and delivering variety ω of good λ from country c to country j . Therefore, firms in country c find it profitable to start production of this variety only if they are the minimum cost suppliers in at least one of the C potential destinations. Conditioning on this, the representative firm maximizes profits by choosing total output and allocating it across the markets it decides to serve; that is, quantities to sell in the domestic and export markets. However, since firms are price takers, in equilibrium prices will be such that every destination served is equally profitable for firms producing a given variety.¹⁷ Otherwise, the firm could reallocate output

across destinations and increase profits. Let us denote this FOB price as $\hat{p}_c(\omega, \lambda)$. The profit maximization problem of a producer of variety ω of good λ in country c can then be written as:

$$\max_h \hat{p}_c(\omega, \lambda) \varepsilon(\omega) A(\lambda, c) h^{\frac{\lambda}{\gamma}} - w_c h - f.$$

Alternatively we can view the firm's problem as one of cost minimization. The combination of a fixed cost and decreasing marginal returns to h implies that firms face a U-shaped average cost curve. The minimum of this curve pins down the firm's employment $h_c^*(\lambda) = f\gamma/[w_c(\lambda - \gamma)]$ and efficient scale $y_c^*(\omega, \lambda) \equiv \varepsilon(\omega)A(\lambda, c)h_c^*(\lambda)^{\frac{\lambda}{\gamma}}$. The minimum average production cost, denoted $m_c^*(\omega, \lambda)$ is:

$$m_c^*(\omega, \lambda) = \kappa(\lambda) \frac{(w_c)^{\frac{\lambda}{\gamma}}}{\varepsilon(\omega)A(\lambda, c)},$$

where $\kappa(\lambda) > 0$ under decreasing marginal returns.¹⁸

Under perfect competition, profit maximizing firms equate output price, $\hat{p}_c(\omega, \lambda)$, to marginal cost. In turn, free entry implies that every firm produces at the efficient scale, where marginal cost equals average cost, $m_c^*(\omega, \lambda)$. Therefore, assuming that there is a continuum of small potential entrants, the industry's average cost function is perfectly elastic at the FOB price $\hat{p}_c(\omega, \lambda)$ that satisfies the zero-profit condition $\hat{p}_c(\omega, \lambda) = m_c^*(\omega, \lambda)$. As a result, the industry's unit cost of producing and delivering variety ω of good λ from country c to country j is:

$$\mu_{cj}(\omega, \lambda) = \tau_{cj} \cdot m_c^*(\omega, \lambda),$$

where $\tau_{cj} \geq 1$ is the iceberg trade cost from c to j . Two remarks are in order. First, we assume that the industry with perfect skill substitutability features no producer heterogeneity ($\varepsilon = 1$) and no transport costs. If there is positive production in this industry in all countries, then these assumptions imply that unit costs in this industry are equalized across countries and so are wage rates. In summary, these assumptions yield wage equalization in equilibria with incomplete specialization. Second, note that unit costs are inversely related to fundamental productivity. Therefore, under Property 1, countries with higher skill dispersion will, ceteris paribus, have relatively lower unit costs of producing varieties in industries with higher substitutability.

5. The pattern of trade

The analysis focuses on equilibria with incomplete specialization in production, in which every country produces positive output in every industry. As explained above, this leads to wage equalization across countries, thus we choose labour as the numeraire and let $w_c = 1$ for all c . As a result, the pattern of bilateral trade is indeterminate in the industry that exhibits perfect skill substitutability. Whenever skills are imperfectly substitutable, however, geographical specialization is determined by the locations of the minimum cost suppliers to each destination, which are functions of trade costs and both fundamental and random productivity shocks. We have seen how the interaction of skill dispersion and skill substitutability generates fundamental productivity differences across industries and countries. This section shows how the pattern of trade is in turn fully determined by fundamental productivity.

At the micro level our framework features worker heterogeneity and non-linear returns to scale; however, at the industry level, it is isomorphic to the perfectly competitive model with homogeneous workers, CRS and cross-country technological differences in Costinot et al. (2012).¹⁹ We therefore follow their derivation of the pattern of trade.

¹⁵ In contrast, BGP uses a grid of λ s defined over the $[0.1, 1]$ interval.

¹⁶ The only exception is one large outlying λ in the parametrization where $\gamma = 0.7$.

¹⁷ In particular, producer prices faced by every firm producing a given variety in the same location will be equalized across destinations.

¹⁸ Specifically, $\kappa(\lambda) \equiv f^{\frac{\lambda}{\gamma-1}}(\lambda - \gamma)^{\frac{\lambda}{\gamma-1}} \gamma^{-\frac{\lambda}{\gamma}} \lambda$.

¹⁹ In particular, production technologies in our model can be equivalently interpreted as depending on h homogeneous workers, with fundamental productivity equal to $A(\lambda, c)$ in industry λ and country c . Furthermore, as shown in the previous section, under free entry, industry unit costs are constant, as with CRS technologies.

As a first step, we provide an expression for bilateral exports in industries with imperfect skill substitutability. For $\lambda \in \Lambda_{-1}$, let $x_{cj}(\lambda) \equiv \sum_{\omega \in \Omega_{cj}(\lambda)} x_j(\omega, \lambda)$ denote the value of total exports from country c to country j in industry λ , where $\Omega_{cj}(\lambda) \equiv \left\{ \omega \in \Omega \mid \mu_{cj}(\omega, \lambda) = \min_{1 \leq c' \leq C} \mu_{c'j}(\omega, \lambda) \right\}$ is the subset of varieties exported by country c to country j in industry λ . Then,

Lemma 2. For any $\lambda \in \Lambda_{-1}$,

$$x_{cj}(\lambda) = \frac{\left[(w_c)^\gamma \tau_{cj}(\lambda) / A(\lambda, c) \right]^{-\theta}}{\sum_{c'=1}^C \left[(w_{c'})^\gamma \tau_{c'j}(\lambda) / A(\lambda, c') \right]^{-\theta}} \alpha(\lambda) E_j. \tag{5}$$

Proof. The result follows from the proof of Lemma 1 in the online appendix of Costinot et al. (2012), using the fact that $\varepsilon(\omega)$ i.i.d. Fréchet implies that $z(\omega, \lambda, c) \equiv \varepsilon(\omega) A(\lambda, c)$ i.i.d. Fréchet with scale parameter $A(\lambda, c)$; i.e. $F(z) = \exp[-(z/A(\lambda, c))^{-\theta}]$, for all $z > 0$. Note that $\kappa(\lambda)$ varies by industry, but not by country, and therefore cancels out in Eq. (5).

Intuitively, in industries with imperfect skill substitutability, the location of the minimum cost supplier of any single variety is indeterminate because it depends on the random component of productivity. However, because these shocks are purely idiosyncratic, a country with a higher fundamental productivity will, ceteris paribus, capture a higher proportion of the industry's varieties imported by consumers in any destination.

In order to isolate the role of skill dispersion in shaping comparative advantage, assume that there are no other exogenous sources of relative cost differences across exporters, i.e. $\tau_{cj}(\lambda) = \tau_{c'j}(\lambda)$ for $\lambda \in \Lambda_{-1}$. Then, under wage equalization, Lemma 2 implies that for any importer j and any pair of exporters c and c' , the ranking of relative fundamental productivities fully determines the ranking of relative exports to a given market j . That is, for any two industries $\lambda, \lambda' \in \Lambda_{-1}$:

$$\frac{A(\lambda, c')}{A(\lambda, c)} \leq \frac{A(\lambda', c')}{A(\lambda', c)} \iff \frac{x_{cj}(\lambda)}{x_{c'j}(\lambda)} \leq \frac{x_{c'j}(\lambda')}{x_{cj}(\lambda')}. \tag{6}$$

In our framework, $A(\lambda, c)$ is log-supermodular in λ and c , when countries are ordered according to increasing skill dispersion. This yields the main result of the paper linking skill dispersion, comparative advantage and trade flows.

Proposition 3. Consider an equilibrium with incomplete specialization in production, in which skill dispersion is the only source of relative cost differences across countries, i.e. $\tau_{cj}(\lambda) = \tau_{c'j}(\lambda)$ for $\lambda \in \Lambda_{-1}$. Then, under Property 1, a country with relatively higher dispersion of skills has a comparative advantage, and therefore exports relatively more to any destination, in sectors with higher substitutability $\lambda \in \Lambda_{-1}$.

Proof. It follows immediately from Property 1 and the ranking of relative exports in Eq. (6).

We now briefly explain how to solve for the remaining endogenous variables of the model. The mass of firms in industries with imperfect skill substitutability, denoted $M_c(\lambda)$, can be determined from the market clearing condition that total expenditure on a given industry's varieties equals the sum of the revenues of domestic producers of these varieties. Under wage equalization, equilibrium firm revenues are $h^*(\lambda) + f$, obtained from Section 4. Therefore, the market clearing condition in industry $\lambda \in \Lambda_{-1}$ can be written as $\sum_j x_{cj}(\lambda) = M_c(\lambda)[h^*(\lambda) + f]$, which, together with the solutions for trade flows in Eq. (5), can be used to solve for $M_c(\lambda)$ as a function of expenditure levels E_c . The allocation of labour in this industry,

denoted $L_c(\lambda)$, satisfies $L_c(\lambda) = M_c(\lambda)[h^*(\lambda) + f]$. In turn, the allocation of labour in the industry with perfect skill substitutability is $L_c(1) = L_c - \sum_{\lambda \in \Lambda_{-1}} L_c(\lambda)$, ensuring full employment.²⁰ From this, we can compute the mass of firms and total exports to the rest of the world, $M_c(1)$ and $\sum_j x_{cj}(1)$, using the solution for $h^*(1)$ and imposing market clearing. Finally, aggregate expenditure levels are obtained by imposing the equality of income and expenditure in every country, so that $E_c = \sum_{\lambda \in \Lambda} \sum_j x_{cj}(\lambda)$ for all $c = 1, \dots, C$.

6. Observable skills

This section extends the model to account for observable skills. The equilibrium features sorting by industry on observable skills. As in the benchmark model, however, unobservable skill dispersion fully determines the pattern of trade in equilibria with wage equalization.

Suppose there are N types of workers, indexed by $n = \{1, \dots, N\}$. The input of a worker i of type n is now the product of observable and unobservable components of ability, denoted q_n and a_{ni} respectively. By definition, every worker of type n is endowed with the same skill q_n , which is observable to both firms and workers.²¹ Component a_{ni} , however, can vary across workers of the same type. As before, we assume a_{ni} is unobservable to firms. The structure of the labour market is therefore unchanged, except that the equilibrium now requires determining N type-specific market clearing wages in each country, denoted $w_{n,c}$.

In this context, the technology in Eq. (2) of a firm in industry λ naturally extends to:

$$y = \varepsilon \left[\sum_{n=1}^N \int_0^{h_n} (q_n a_{ni})^\lambda di \right]^\gamma, \tag{7}$$

where h_n is the number of workers of type n employed.²² Note that if $N = q_N = 1$, we obtain Eq. (2).

Let $G_{n,c}(a)$ be the distribution of unobserved skills of type n in country c . To keep matters as close to the benchmark model as possible, we assume that unobservable skills are identically distributed across skill types, i.e. $G_{n,c}(a) = G_c(a)$ for every n , as in Grossman and Maggi (2000).²³ Without loss of generality, we further restrict the analysis to two observable types $n = \{1, 2\}$, with $q_1 < q_2$ and refer to type-1 and type-2 workers as unskilled and skilled, respectively. Together with random matching on unobservable skills, these assumptions allow us to express output as $y = \varepsilon A(\lambda, c)[h_1 q_1^\lambda + h_2 q_2^\lambda]^\gamma$, where $A(\lambda, c) \equiv (\int \alpha^\lambda dG_c(a))^\gamma$ is defined exactly as in Section 3.

Dropping the country and industry indices momentarily in order to lighten notation, the firm chooses h_1 and h_2 in order to maximize $\hat{p}y - w_1 h_1 - w_2 h_2 - f$. The first-order condition for $h_n, n = \{1, 2\}$, is $\hat{p} \varepsilon A [h_1 q_1^\lambda + h_2 q_2^\lambda]^{\gamma/\lambda - 1} q_n^\lambda (\gamma/\lambda) \leq w_n$, with equality if $h_n > 0$.²⁴ Note that

²⁰ Incomplete specialization requires $L_c(1) > 0$ for all c , which can be ensured by appropriate choice of labour endowments and expenditure shares across countries.

²¹ One possible interpretation is that n indexes education–experience cells and q_n is the expected value of the (log) ability of workers belonging to cell n . This interpretation would require the normalization $E_n[\log(a_{ni})] = 0$ for all n .

²² Recall that in the benchmark model without observable skills, there are no skill-specific tasks, i.e. total output is simply the integral of individual inputs a_i^λ across employees, regardless of a_i . By applying the same concept in the present context, we obtain Eq. (7).

²³ Allowing for 'heteroskedasticity' in unobservable skills across countries implies that unobserved skill dispersion for each type of worker operates as an independent source of comparative advantage. Although potentially relevant in an empirical analysis, this extension would make the analysis cumbersome, without adding further insight.

²⁴ The production function is quasiconcave in (h_1, h_2) , since we still assume $0 < \gamma < \lambda$. The first-order conditions are thus necessary and sufficient for a maximum.

both first-order conditions hold with equality only in a threshold industry λ_c^* , satisfying

$$\left(\frac{q_2}{q_1}\right)^{\lambda_c^*} = \frac{w_{2,c}}{w_{1,c}} \tag{8}$$

Producers in industry λ_c^* are indifferent between both types of workers so, for expositional purposes, we focus on equilibria such that industry λ_c^* hires skilled workers.²⁵ Since $q_1 < q_2$, firms in industries $\lambda \geq \lambda_c^*$ exclusively employ skilled workers, while firms in industries $\lambda < \lambda_c^*$ exclusively employ unskilled workers. Therefore, the equilibrium displays sorting on observable skills and random matching on unobservable skills.²⁶ Although the present model considers only two factors, the result is reminiscent of the sorting pattern obtained in Costinot and Vogel (2010) with a continuum of sectors and factors. The key similarity lies in the assumptions of log-supermodularity of the sector-skill specific productivity in their paper and of q^λ in this paper.

Since firms hire a single type of workers, the solutions to the profit maximization problem are very similar to the benchmark model. In particular, under free-entry, it is straightforward to show that the industry's unit cost of producing and delivering variety ω from country c to country j , is $\mu_{cj}(\omega, \lambda) = \tau_{cj}(\lambda) \kappa_{n\lambda,c} (w_{n\lambda,c})^{\gamma/\lambda} / [\varepsilon(w) A(\lambda, c)]$, where subscript n_λ indicates the type of workers (optimally) employed in industry λ ; i.e. $n_\lambda = 1$ if $\lambda < \lambda_c^*$ and $n_\lambda = 2$ otherwise. In turn, $\kappa_{n\lambda,c} \equiv f^{\frac{\lambda-\gamma}{\lambda}} (\lambda-\gamma)^{\lambda-1} \gamma^{-\lambda} \lambda (q_{m\lambda})^{-\gamma}$ is a function of parameters that vary across industries and, possibly, countries.²⁷ Having determined unit costs, total bilateral exports at the industry level are derived as in Lemma 2:

$$x_{cj}(\lambda) = \frac{\left[\kappa_{n\lambda,c} (w_{n\lambda,c})^{\frac{\gamma}{\lambda}} \tau_{cj}(\lambda) / A(\lambda, c)\right]^{-\theta}}{\sum_{c'=1}^C \left[\kappa_{n\lambda,c'} (w_{n\lambda,c'})^{\frac{\gamma}{\lambda}} \tau_{c'}(\lambda) / (A(\lambda, c'))\right]^{-\theta}} \alpha(\lambda) E_j \tag{9}$$

As in the benchmark model, we can now establish the pattern of trade for the class of equilibria featuring wage equalization across countries. In this case, $w_{n,c} = w_n$ for all c together with Eq. (8) imply that the assignment of observable skills to industries is the same in every country; i.e. $\lambda_c^* = \lambda^*$ for all c . In the absence of other sources of comparative advantage, it is easy to verify that Eq. (6) holds and thus relative trade flows are fully determined by unobservable skill dispersion. The next result immediately follows from these observations, along the lines of Proposition 3.

Proposition 4. Consider an equilibrium with wage equalization across countries in the model with multiple observable skills. Assume that transport costs are not a source of relative cost differences across countries, i.e. $\tau_{cj}(\lambda) = \tau_{c'j}(\lambda)$ for $\lambda \in \Lambda$. Then, under Property 1, a country with relatively higher dispersion of unobservable skills has a comparative advantage, and therefore exports relatively more to any destination, in sectors with higher substitutability.

We note that wage equalization can be ensured in a similar way as in the benchmark model and, more generally, as in standard Heckscher–Ohlin theory. That is, by an appropriate choice of endowments and expenditure shares such that every country is incompletely specialized in a set of N , freely traded, homogeneous goods, each employing a different type of workers.

Finally, we would like to remark that in our model, because of the assumption of wage equalization across countries, observable skill dispersion is not a source of comparative advantage among differentiated goods.²⁸ However, the analysis in this section is not intended to provide definitive conclusions regarding the effect of observable skill dispersion on trade flows, but rather to show that observable skills can be incorporated into our model in a natural way, without necessarily changing our core results regarding unobservable skill dispersion.

7. Quantitative analysis

Next, we return to the model outlined in Sections 2–5 to perform a quantitative assessment of the effect of skill dispersion on trade flows. For this exercise we consider a set of 63 manufacturing industries and 18 countries for which we can observe both bilateral trade flows and the distribution of workers' skills.²⁹

7.1. Eliminating cross-country differences in skill dispersion

We undertake a set of counterfactual experiments aimed at inferring how changes in comparative advantage due to skill dispersion alter trade flows. To do so we fix a reference country and study the general equilibrium responses of its trade flows with the rest of the world (ROW) after removing all comparative advantage due to skill dispersion. More specifically, after choosing a reference country 0, we set each industry's $A(\lambda, c) = A(\lambda, 0)$ for all the other countries c ; then we solve for the counterfactual equilibrium and we compare both exports and imports of country 0 to the ROW, industry by industry, against those in the benchmark equilibrium for which the $A(\lambda, c)$ is constructed from data.³⁰ Consistent with the model, and to minimize confounding effects, we perform the quantitative analysis in a setting in which (i) wages are equalized in equilibrium, (ii) consumers have identical preferences, i.e. $a(\lambda, c) = a(\lambda)$ for all c , and (iii) skill dispersion is the only source of comparative advantage. The third assumption requires $\tau_{cj}(\lambda) = \tau_{c'j}(\lambda)$ for every c, j and λ . Below we describe all the experiments, then we discuss their implementation.

Change in exports, by industry. Under the assumptions (i)–(iii) above, we use Eq. (5) to obtain an expression for the exports of country 0 to the ROW in industry λ . Counterfactual equilibrium exports, relative to those in the benchmark equilibrium, are

$$\frac{\sum_j x'_{0j}(\lambda)}{\sum_j x_{0j}(\lambda)} = \frac{\sum_j \frac{\tau'_{0j}}{\tau_{0j}} E'_j}{\sum_j \frac{[\tau_{0j}/A(\lambda, 0)]^{-\theta}}{[\tau_{c'j}/A(\lambda, c')]^{-\theta}} E_j} \tag{10}$$

where primes denote counterfactual variables. For example, $x_{0j}(\lambda)$ and $x'_{0j}(\lambda)$ are exports from 0 to j in industry λ in, respectively, the benchmark and counterfactual equilibria. Further below we explain how we compute Eq. (10) using data and, then, we use it to assess the impact of skill dispersion on each industry's exports.

²⁵ We assume $\lambda_c^* \in \Lambda$. This will be the case when, for example, Λ is an interval in the real line.

²⁶ Intuitively, sorting is driven by the assumption of no type-specific tasks embedded in Eq. (7) together with the fact that returns to individual skills increase in λ .

²⁷ $\kappa_{n\lambda,c}$ will vary across countries as long as $q_{n\lambda}$ does. This can occur when the threshold industry λ_c^* differs across countries so that, in a given industry λ , countries employ workers of different types.

²⁸ Standard Heckscher–Ohlin comparative advantage may emerge in the homogeneous good sectors, in which skills are perfectly substitutable. We do not focus on this well established result.

²⁹ Our sample of industries and countries is identical to the one studied in BGP, with the exclusion of Chile. We are unable to measure Chilean domestic absorption rates from OECD-STAN data and thus we drop this country from the sample (see Online Appendix).

³⁰ These counterfactual experiments are similar to those presented in Costinot et al. (2012) to study Ricardian comparative advantage, with two important differences: (i) we allow for counterfactual changes in aggregate expenditures and, (ii) we also report results at a more disaggregated (industry) level.

It is also interesting to study how this ratio changes across industries. Consider any two industries λ_1 and λ_2 ; then it is easy to show that

$$\frac{\sum_j x'_{oj}(\lambda_1)}{\sum_j x_{oj}(\lambda_1)} \leq \frac{\sum_j x'_{oj}(\lambda_2)}{\sum_j x_{oj}(\lambda_2)} \tag{11}$$

if and only if

$$\sum_j \frac{[\tau_{oj}/A(\lambda_2, 0)]^{-\theta}}{\sum_{c'} [\tau_{c'j}/A(\lambda_2, c')]^{-\theta}} E_j \leq \sum_j \frac{[\tau_{oj}/A(\lambda_1, 0)]^{-\theta}}{\sum_{c'} [\tau_{c'j}/A(\lambda_1, c')]^{-\theta}} E_j \tag{12}$$

The last expression suggests that, following a loss of comparative advantage, the decline of country 0's relative exports is larger in industries in which this country had a relatively higher fundamental productivity advantage with respect to the ROW. This can be seen most clearly when trade costs are symmetric, i.e. $\tau_{c'j} = \tau_{oj}$ for all c' and j . In this case Eq. (12) simplifies to

$$\frac{A(\lambda_2, 0)^\theta}{\sum_{c'} A(\lambda_2, c')^\theta} \leq \frac{A(\lambda_1, 0)^\theta}{\sum_{c'} A(\lambda_1, c')^\theta} \tag{13}$$

If the above inequality holds, country 0 has a comparative advantage in industry λ_1 relative to the ROW and, using Eq. (11), its counterfactual exports to the ROW should drop relatively more in industry λ_1 . For example, if we take a high skill dispersion country (e.g. the US) as country 0, its fundamental productivities in the initial equilibrium would be relatively larger (relative to the ROW) in industries with higher λ . It follows that, when plotting $\sum_j x'_{oj}(\lambda)/\sum_j x_{oj}(\lambda)$ as a function of λ , we expect to observe, on average, a negatively sloping curve. The steepness of this curve indicates the strength of CA generated by skill dispersion. In contrast, if the reference country has low skill dispersion (e.g. Germany), that slope is expected to be positive. In the numerical experiments below we do not assume symmetric trade costs and, therefore, we also allow for some confounding effects.

Change in imports, by industry. Similarly, the change in imports to country 0 from the ROW in the counterfactual equilibrium can be expressed as

$$\frac{\sum_{c \neq 0} x'_{c0}(\lambda)}{\sum_{c \neq 0} x_{c0}(\lambda)} = \frac{1 + \frac{[\tau_{00}(\lambda)/A(\lambda, 0)]^{-\theta}}{\sum_{c \neq 0} [\tau_{c0}(\lambda)/A(\lambda, c)]^{-\theta}}}{1 + \frac{\tau_{00}(\lambda)^{-\theta}}{\sum_{c \neq 0} \tau_{c0}(\lambda)^{-\theta}}} \tag{14}$$

Comparing such changes in any two industries λ_1 and λ_2 under the hypothesis that trade costs do not generate comparative advantage, one obtains

$$\frac{\sum_c x'_{c0}(\lambda_1)}{\sum_c x_{c0}(\lambda_1)} \leq \frac{\sum_c x'_{c0}(\lambda_2)}{\sum_c x_{c0}(\lambda_2)}$$

if and only if

$$\frac{A(\lambda_1, 0)^\theta}{\sum_{c \neq 0} [\tau_{c0}/A(\lambda_1, c)]^{-\theta}} \leq \frac{A(\lambda_2, 0)^\theta}{\sum_{c \neq 0} [\tau_{c0}/A(\lambda_2, c)]^{-\theta}} \tag{15}$$

This condition is similar to Eq. (13), with the difference due to the presence of trade costs. Simply put, the condition suggests that if we take a high skill dispersion country (US) as country 0, we expect its fundamental productivities to be relatively larger (compared to the ROW) in industries with higher λ . Eq. (15) suggests that, when plotting

$\sum_c x_{c0}/\sum_c x_{c0}(\lambda)$ as a function of λ , we should observe a positive slope. This is quite intuitive since a positive slope means that, when the US loses its comparative advantage in high substitutability industries, the ROW exports relatively more to the US in those industries.³¹ Like before, condition in Eq. (15) is informative about the slope of relative changes in imports across industries, while the relative change within any industry is summarized by Eq. (14).

7.2. Calibration

To perform the quantitative exercises described above one needs to assign values to a subset of model parameters. Here we overview these parameters. Estimates and further details are deferred to the Online Appendix. Once parameter values are obtained, we solve for bilateral trade flows and aggregate expenditures in both the benchmark and counterfactual equilibria, as outlined at the end of Section 5. This allows us to compute the numerical counterparts of expressions (10) and (14) for our sample of countries and industries.

7.2.1. Technology

The crucial technology parameters in Eq. (2) are the industry-specific λ s, which determine the elasticity of substitution among workers' skills, and γ , which controls the degree of diminishing returns to labour, for a given λ .³² From the analysis in Section 4, it is straightforward to check that the wage bill as a fraction of firm revenue is equal to γ/λ , in industry λ .³³ We therefore use industry-level data to compute the share of labour in value added to approximate γ/λ . The source of the data is the 2002 Economic Census and the variables employed are total compensation³⁴ and value added at the NAICS 6-digit level, aggregated to the 63 industries that are constructed in BGP.

For a given γ , the procedure described above allows us to pin down λ in each of the 63 industries considered. However, implementing our quantitative exercises still requires assigning a value for γ . To gauge the robustness of our results, we parametrize the technology under alternative values of γ and then simply invert the estimated γ/λ ratios to derive the associated sets of λ s. In the context of our model, it is reasonable to impose a value of 1 as an upper-bound on the range of γ . Otherwise, the calibrated lambdas would be larger than 1 in every industry, a highly implausible scenario in which production technologies in every manufacturing industry have concave isoquants.³⁵ Therefore, we posit that $\gamma \in (0, 1)$ and experiment with three values. Namely, we consider a benchmark parameterization of $\gamma = 0.5$ (the middle value of the unit range) but we also report results for γ equal to 0.3 and 0.7.³⁶ More details about the technology parameterization are presented in the Online Appendix.

³¹ In Eq. (15) the transport costs τ generate a different weighting of the A productivities for ROW and this may, in principle, weaken somewhat the result.

³² The empirical literature has so far paid little attention to the estimation of individual skills' substitutability in production. Most studies of heterogeneous labour demand have instead focused on coarser classifications, such young vs old, male vs female, college vs no college. A discussion of relevant issues can be found in the original survey by Hamermesh (1986).

³³ More specifically, for any variety ω of good λ , λ , $(w_c h_c^*(\lambda))/(\bar{p}_c(w, \lambda) y_c^*(w, \lambda)) = \gamma/\lambda$.
³⁴ Results are unaffected when employing total payroll as a measure of wage bill.

³⁵ The ONET database studied in BGP reveals that scores of indicators of low skill substitutability, such as the importance of teamwork, are high in many US manufacturing industries.

³⁶ In a different context, when aggregating across 67 different occupations, Hsieh et al. (2013) acknowledge the lack of information about the elasticity of substitution and simply choose a value of 3, doing robustness checks around this benchmark value. We take an approach similar to theirs. While their setup is different from ours, their assumed elasticity is well within the range we allow for in our application.

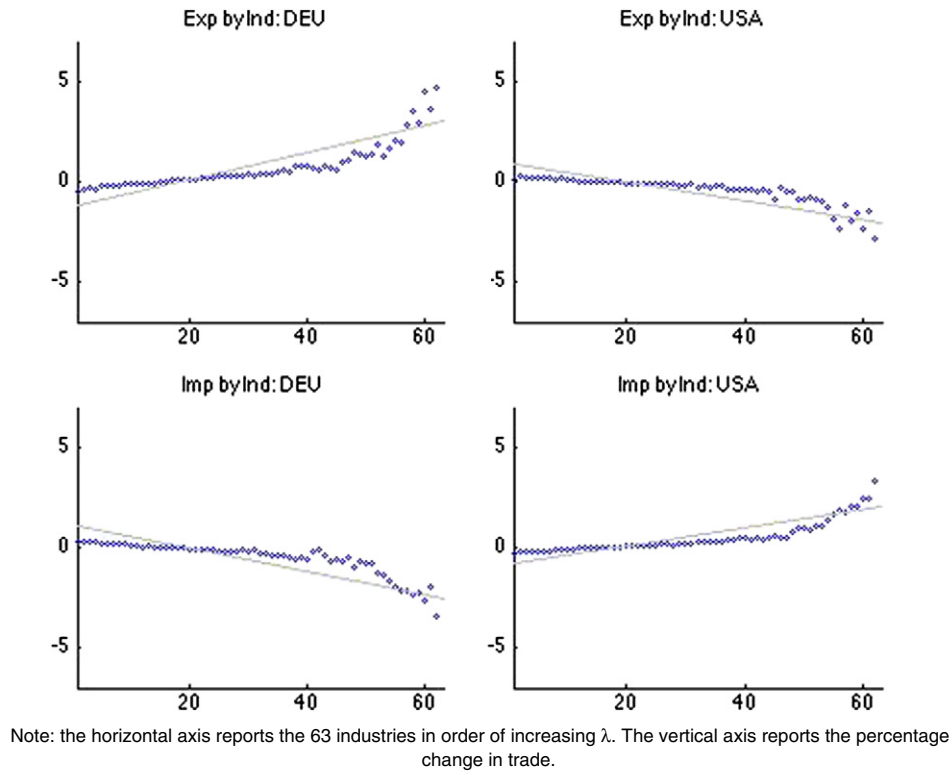


Fig. 2. Changes in trade flows for Germany and US: $\gamma = 0.5$.

7.2.2. Skill distributions

To approximate a country's skill distribution we use the distribution of IALS scores (see BGP for a discussion). We do this for 18 countries.³⁷ Together with the estimated lambdas, we use Eq. (3) to compute fundamental productivity measures $A(\lambda, c)$ for each country c in industry λ .

7.2.3. Transportation costs

The transportation costs $\tau_{cj}(\lambda)$ have an important role in the numerical analysis, capturing all 'residual' sources of unmodelled heterogeneity that affect trade between each country pair (c, j) in industry λ . We estimate $\tau_{cj}(\lambda)$ in two steps: first, we use the OECD-STAN bilateral trade data at the industry level. The $\tau_{cj}(\lambda)$ is chosen to replicate bilateral trade flows by industry; this results in a set of $18 \times 17 \times 63$ elements corresponding to all country–industry triplets. In the Online Appendix we describe how these trade costs are derived as a function of bilateral trade flows in each industry. Next, we decompose (through an OLS projection) each $\tau_{cj}(\lambda)$ into two multiplicative components: (1) a country-pair effect; (2) a country–industry component. These are the two terms $\hat{\tau}_{cj}$ and $\hat{\tau}_j(\lambda)$, which we then use in the simulations. The decomposition $\hat{\tau}_{cj}(\lambda) = \hat{\tau}_{cj} \hat{\tau}_j(\lambda)$ guarantees that transport costs are not a source of comparative advantage.³⁸

7.2.4. Idiosyncratic variety shocks

Parameter θ , characterizing the Fréchet distribution of the Ricardian productivity shocks $\varepsilon(\omega)$, is set to 6.53, following Costinot et al. (2012).

7.2.5. Preferences

To solve for an equilibrium we only need to set the expenditure share parameters $a(\lambda)$, which we assume identical in all countries. Details of how we calibrate these consumption shares are in the Online Appendix.

8. Results

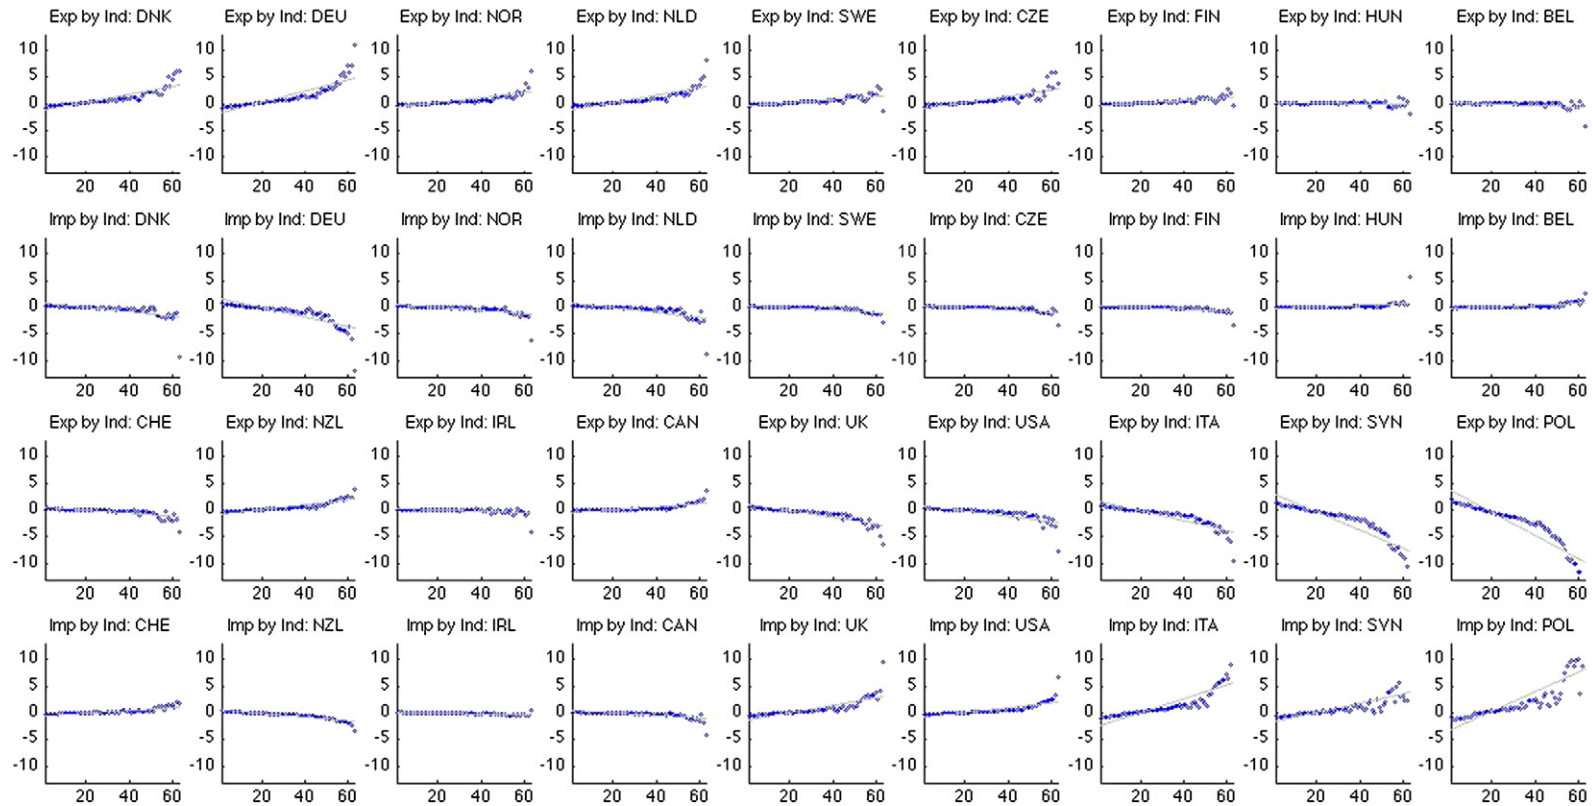
Shutting down differences in relative productivities guarantees the absence of comparative advantage due to skill dispersion. However, the choice of a reference country is important in these experiments because it affects the level of fundamental productivities' in different industries. For this reason we compute a battery of 18 counterfactual experiments, allowing for a different reference country in each one of them. In what follows we plot graphs based on technology estimates obtained using the wage shares measured through total compensation and the benchmark parametrization of $\gamma = 0.5$. In the Online Appendix we report results for alternative values of γ . Fig. 2 plots the percentage changes in imports and exports, setting Germany and the US as reference countries.³⁹ While both countries are at similar stages of development, they are at opposite ends of the skill dispersion range. Each square reports the proportional (percentage) change in exports or imports for the reference country, plotted over the 63 differentiated goods' sectors: industries on the horizontal axis are ordered in terms of increasing λ (i.e. skills' substitutability in production). In each graph we overlay a linear fit, to emphasize the direction and steepness of the changes as the skill substitutability parameter λ becomes larger.

Note: the horizontal axis reports the 63 industries in order of increasing λ . The vertical axis reports the percentage change in trade.

³⁷ The results presented here are produced using 'residual' IALS obtained by regressing IALS scores on individual-level observable characteristics, as documented in BGP. We also perform the same experiments using raw IALS scores, and the results are qualitatively and quantitatively similar. In line with the fact that most (roughly 2/3) of IALS variation comes from the residual component, the experiments with raw scores exhibit effects on trade flows which are between 15 and 30% stronger. More details are available from the authors.

³⁸ Using this decomposition of the transport costs means that the simulations do not exactly match bilateral trade flows in each industry.

³⁹ Changes are computed like in Eqs. (10) and (14). Exports of a country to itself are not included.



Note: the horizontal axis reports the 63 industries in order of increasing λ . The vertical axis reports the percentage change in exports or imports.

Fig. 3. Changes in trade flows across industries: $\gamma = 0.5$.

These plots reveal that the expected patterns of change in exports and imports are borne out by the experiments. Export changes for the highest skill dispersion country (the US) show a decreasing pattern as we move towards higher λ industries, while the slope of the same curve for Germany exhibits a positive slope. The opposite is true for imports, confirming that the loss of comparative advantage is reflected more strongly in the sectors which most benefit from a country's skill dispersion in the benchmark equilibrium.

This quantitative result is further illustrated by Fig. 3 which reports trade changes for the full set of 18 reference countries. The first two rows show, respectively, proportional changes in exports and imports by industry for the nine countries with the lowest skill dispersion. The last two rows do the same for the nine countries with the highest skill dispersion. In each row skill dispersion is increasing as countries change from left to right. From these plots it is apparent that, when we equalize the $A(\lambda, c)$ using reference countries at the lower end of the skill dispersion range (like Denmark and other Nordic countries) export changes become larger as we move right towards industries with higher λ , while the opposite is true for import changes. In contrast, when the reference country in the counterfactual experiments have high skill dispersion like the US, the UK or Italy, changes in trade flows by industry have slopes which are the reversed, as these countries lose their comparative advantage in high λ industries. Interestingly, for countries in the middle of the skill dispersion range (e.g., Hungary, Belgium, Switzerland) there is little or no slope in the curves plotting export and import changes by industry: these are countries where comparative advantage due to skill dispersion is generally weaker in the benchmark equilibrium. Beyond the evolution of trade flow changes across industries, these experiments can also shed some light on the relative intensity of the changes. To this purpose we compute for each country the average of the absolute percentage changes in imports and exports. Table 1 reports these average absolute changes. The table also reports the (raw, not absolute) percentage change experienced in the industries corresponding to the 10th and 90th percentiles in the distribution of estimated λ s. The latter values clearly illustrate the range of variation in percentage changes.

These results confirm that the largest changes are experienced by countries which lie at the opposite ends of the skill dispersion range. Changes are also more pronounced for sectors at the extremes of the λ range, in particular for the highest values of λ . Moreover, looking at the changes in industries at the extremes of the substitutability scale it is apparent that the signs of changes are inverted for countries which lie at different ends of the skill dispersion range. Finally, to summarize the individual country results, we take a weighted average of the absolute changes in exports (imports) across countries, where weights correspond to export (import) shares in the benchmark. This results in an average change of 1.63% in exports and 1.35% in imports. These findings are in line with our theoretical analysis and suggest a non-trivial role for skill dispersion in determining trade flows. In the Online Appendix we also report results for simulations in which we solve and perturb the benchmark under different technology parameters, in particular γ . Setting γ to, respectively, 0.3 and 0.7, the weighted average absolute change in exports are, respectively, 0.85% and 4.50%. To put these results in perspective, one would obtain a 5.68% average change in absolute value of exports when using counterfactual results presented by Costinot et al. (2012) from a similar experiment designed to quantitatively assess Ricardian comparative advantage.⁴⁰

⁴⁰ This number is obtained by taking an un-weighted average of the absolute value of export changes presented in column (1) of Table 7 of Costinot et al. (2012). Some caution should be exercised in comparing results since several details of the counterfactual experiments differ across papers, including the samples of countries considered.

Table 1
Changes in trade flows ($\gamma = 0.5$).

Country	Exports			Imports		
	(1)	(2)	(3)	(4)	(5)	(6)
	Average	10%	90%	Average	10%	90%
DNK	0.93	-0.25	3.81	0.50	0.12	-1.08
NLD	0.92	-0.23	2.71	0.60	0.16	-1.92
NOR	0.59	-0.17	1.30	0.37	0.10	-1.21
FIN	0.48	-0.13	1.70	0.37	0.09	-1.04
DEU	1.03	-0.24	3.54	0.82	0.22	-2.34
SWE	0.46	-0.09	1.18	0.20	0.07	-0.60
CZE	0.56	-0.04	3.32	0.17	0.04	-0.44
BEL	0.15	-0.01	0.18	0.09	-0.02	0.32
HUN	0.22	0.01	1.54	0.04	-0.01	0.10
CHE	0.20	0.04	-0.54	0.15	-0.06	0.59
CAN	0.55	-0.18	1.85	0.40	0.18	-1.84
NZL	0.31	-0.04	0.48	0.26	0.02	-0.47
IRL	0.28	0.01	-0.41	0.08	0.02	-0.06
UK	0.54	0.22	-1.82	0.57	-0.20	2.05
USA	0.67	0.16	-1.96	0.68	-0.18	2.08
ITA	0.97	0.26	-2.00	1.44	-0.31	3.68
SVN	1.67	0.42	-4.15	0.84	-0.35	3.29
POL	2.71	0.75	-7.45	2.34	-0.73	6.46

Notes:

- Columns (1) and (3) report the average of absolute percentage changes in trade; columns (2), (4), (5) and (6) report raw (not absolute) percentage changes in trade for the industries corresponding to the 10th and 90th percentile in the distribution of estimated λ 's (skill substitutability).
- The countries are ranked by skill dispersion.
- 10th percentile and 90th percentile are the 6th and 58th industries ranked by λ 's.

Finally, to gain a sense of the welfare implications of these counterfactuals we also computed changes in the price index for the set of differentiated products. Under the assumption that $\gamma = 0.5$, we focused on counterfactuals involving only the highest and lowest skill dispersion countries (respectively, Denmark and Poland). This exercise suggests that shutting down cross-country differences in skill-dispersion would increase the price index by 0.18% when taking Denmark as the reference country. The same exercise results in an average decrease in the price index of 0.35% when Poland is the reference country. These results combine the effect of removing differences in skill dispersion on comparative advantage and on the absolute level of productivity. When $\gamma = 0.5$, many industries feature calibrated λ 's above 1, which explains why raising skill dispersion to the level of Poland has a beneficial effect on all countries' price indices. This effect more than offsets the increase in the price index due to the elimination of comparative advantage when a high fraction of λ s is above one. The opposite happens when $\gamma = 0.3$ and most λ s are below one (implying convex isoquants): the Denmark experiment yields an average differentiated product price index decrease of 0.03% for the Denmark case and an increase of 0.065% for the Poland case.

9. Conclusions

Relative differences in the distribution of factors of production are central to the classical theory of international trade. The Heckscher–Ohlin–Samuelson factor proportion model stresses the idea that cross-country differences in aggregate factor endowments play a major role in predicting trade flows. This paper belongs to a line of research that emphasizes that the entire distribution of a productive factor can enhance our understanding of trade patterns.

We develop a theoretical framework where, because skills are unobservable by firms, workers and firms are randomly matched. Skill dispersion affects industries differently because some technologies are more capable of substituting skills across production tasks than others. All industries in each country inherit the population distribution of

unobservable skills and, as a result, firms in sectors with lower substitutability are relatively more productive in countries with lower skill dispersion. The paper therefore shows how differences in the dispersion of human capital inputs may lead to ‘Ricardian-looking’ differences in measured labour productivity at the country–industry level. In Ricardian models such productivity differences are often assumed to be the result of access to different (broadly defined) technologies. Our findings suggest that productivity wedges may also arise as the by-product of cross-country differences in the distribution of skills, when the latter are not properly accounted for. Higher productivity implies that countries with low unobservable skill dispersion specialize in (and export) goods produced under low skill substitutability. This result survives the introduction of cross-country differences in observable skills and provides a theoretical underpinning to the empirical analysis in BGP.

Relative to the seminal contribution of Grossman and Maggi (2000) we adopt a specific production function and we are therefore able to specify theoretical conditions under which, even in the presence of convex isoquants in all industries, comparative advantage and the pattern of trade are well defined.

In a calibration exercise we show that productivity wedges associated to skill dispersion are quantitatively important determinants of trade flows. We find that eliminating all skill dispersion differences produces a reallocation of production and an average absolute change in exports of almost 2%. Moreover these changes exhibit substantial heterogeneity across countries and sectors.

Appendix A

To simplify notation we define $B(\lambda, c)$ such that $B\gamma(\lambda, c) = A(\lambda, c)$. It is immediate to show that $B(\lambda, c)$ is log-supermodular if and only if $A(\lambda, c)$ is log-supermodular. Proof is conducted in terms of $B(\lambda, c)$ for notational simplicity.

A.1. Proof of proposition 1

By definition of log-supermodularity we need to prove that, if $G_c(a)$ is a mean-preserving spread of $G_c(a)$, then:

$$\frac{\partial \log B(\lambda, c)}{\partial \lambda} \leq \frac{\partial \log B(\lambda, c')}{\partial \lambda}.$$

The partial derivative has the following expression:

$$\frac{\partial \log B(\lambda, c)}{\partial \lambda} = \frac{1}{\lambda} \int a^\lambda \log(a) dG_c(a) - \frac{1}{\lambda^2} \log\left(\int a^\lambda dG_c(a)\right). \tag{A-1}$$

A mean-preserving spread of $g(a, c)$ increases the second term of the right-hand side of (A-1) by definition, since a^λ is a concave function. A sufficient condition for the first term of (A-1) to increase with a mean-preserving spread in $g(a, c)$ is that $k(a) = a^\lambda \log a$ is a convex function which is verified if its second derivative with respect to a is positive for every value of a . i.e. $\log a < \frac{2\lambda-1}{(1-\lambda)\lambda}$. Note that the right-hand side of this inequality is continuous and increasing in λ , and $\lim_{\lambda \rightarrow 1} \frac{2\lambda-1}{(1-\lambda)\lambda} = \infty$. Then, if a is bounded above by a_{\max} , there exists a value $\lambda_{\min} < 1$ such that $\log a_{\max} = \frac{2\lambda_{\min}-1}{(1-\lambda_{\min})\lambda_{\min}}$. If $\lambda > \lambda_{\min}$ then $\frac{\partial \log B(\lambda, c)}{\partial \lambda}$ increases with a mean preserving spread of $g(a, c)$.

We conclude that given an upper bound on the support of the skill distribution a_{\max} , we can find a value λ_{\min} satisfying $0 < \lambda_{\min} < 1$ such that $A(\lambda, c)$ is log-supermodular in λ and c , for $\lambda_{\min} \leq \lambda \leq 1$. This completes the proof of Proposition 1.

A.2. Property 1 when skill distributions are log-normal or Pareto

- (i) *Pareto distribution* – Under the assumption that skills follow a Pareto distribution with mean μ and standard deviation σ , B takes the following expression:⁴¹

$$B = \frac{\mu^2 + \sigma^2 - \sigma\sqrt{\mu^2 + \sigma^2}}{\mu} \left(\frac{\sigma + \sqrt{\mu^2 + \sigma^2}}{\sigma + \sqrt{\mu^2 + \sigma^2} - \lambda\sigma} \right)^{\frac{1}{\lambda}}$$

Since B is twice differentiable in σ and λ , the result in Proposition 4 is equivalent to B being log-supermodular in λ and σ , that is $\frac{\partial^2 \log B}{\partial \sigma \partial \lambda} > 0$. The expression for the cross partial derivative is the following:

$$\frac{\partial^2 \log B}{\partial \sigma \partial \lambda} = \frac{\sigma(\sqrt{\mu^2 + \sigma^2} - \sigma)}{\sqrt{\sqrt{\mu^2 + \sigma^2}[\sigma(1-\lambda) + \sqrt{\mu^2 + \sigma^2}]}}. \tag{A-2}$$

$\lambda < 1$ ensures B is log-supermodular in λ and σ .

- (ii) *Log-normal distribution* – If the distribution of skills a is lognormal on the support $[0, \infty)$ with mean μ and standard deviation σ then B takes the following form:

$$B = e^{\log \mu - \frac{1}{2}\lambda \log\left(\frac{\sigma^2}{\mu^2} + 1\right)}.$$

It is easy to show that under this distribution, B is log-supermodular for every $\lambda > 0$ since the following expression is always positive:

$$\frac{\partial^2 \log B}{\partial \sigma \partial \lambda} = \frac{\sigma}{\mu^2 + \sigma^2}.$$

Appendix B. Supplementary data

Supplementary data to this article can be found online at <http://dx.doi.org/10.1016/j.jinteco.2013.11.005>.

References

Akerlof, G.A., 1970. The market for “lemons”: quality uncertainty and the market mechanism. *Q. J. Econ.* 488–500.
 Benabou, R., 2005. Inequality, technology and the social contract. *Handb. Econ. Growth* 1, 1595–1638.
 Bombardini, M., Gallipoli, G., Pupato, G., 2012a. How robust is the skill-dispersion-complementarity hypothesis? Mimeo, The University of British Columbia.
 Bombardini, M., Gallipoli, G., Pupato, G., 2012a. Skill dispersion and trade flows. *Am. Econ. Rev.* 102 (5), 2327–2348.
 Bougheas, S., Riezman, R., 2007. Trade and the distribution of human capital. *J. Int. Econ.* 73 (2), 421–433.
 Costinot, A., Donaldson, D., Komunjer, I., 2012. What goods do countries trade? a quantitative exploration of Ricardo’s ideas. *Rev. Econ. Stud.* 79 (2), 581–608.
 Costinot, A., Vogel, J., 2010. Matching and inequality in the world economy. *J. Polit. Econ.* 118 (4), 747–786.
 Grossman, G.M., 2004. The distribution of talent and the pattern and consequences of international trade. *J. Polit. Econ.* 112 (1), 209–239.
 Grossman, G.M., Maggi, G., 2000. Diversity and trade. *Am. Econ. Rev.* 1255–1275.

⁴¹ The Pareto distribution is characterized by a shape parameter k and location parameter a_{\min} , i.e. the cumulative distribution of ability is given by $G(a) = 1 - (\frac{a_{\min}}{a})^k$ with $a_{\min} > 0$ and $k > 2$. We could have written B as a function of those parameters: $B = a_{\min} (\frac{k}{k-\lambda})^{\frac{1}{\lambda}}$. Since we are interested in a mean-preserving increase in variance, we express the B as a function of μ and σ , which are related to shape and location parameters according to the following equations:

$$a_{\min} = \frac{\mu^2 + \sigma^2 - \sigma\sqrt{\mu^2 + \sigma^2}}{\mu}, \quad k = \frac{\sigma + \sqrt{\mu^2 + \sigma^2}}{\sigma}.$$

- Hamermesh, D.S., 1986. The demand for labor in the long run. *Handb. Labor Econ.* 1, 429–471.
- Helpman, E., Itskhoki, O., Redding, S., 2010. Inequality and unemployment in a global economy. *Econometrica* 78 (4), 1239–1283.
- Hsieh, C.-T., Hurst, E., Jones, C.I., Klenow, P.J., 2013. The allocation of talent and us economic growth. Tech. rep. National Bureau of Economic Research.
- Kremer, M., 1993. The O-ring theory of economic development. *Q. J. Econ.* 551–575.
- Ohnsorge, F., Trefler, D., 2007. Sorting it out: international trade with heterogeneous workers. *J. Polit. Econ.* 115 (5), 868–892.
- Romalis, J., 2004. Factor proportions and the structure of commodity trade. *Am. Econ. Rev.* 67–97.
- Ross, S.A., 1981. Some stronger measures of risk aversion in the small and the large with applications. *Econometrica J. Econ. Soc.* 621–638.