The Evolution of Comparative Advantage: Measurement and Welfare Implications*

Andrei A. Levchenko
University of Michigan
and NBER

Jing Zhang
University of Michigan

August 3, 2011

Abstract

Using an industry-level dataset of production and trade spanning 75 countries and 5 decades, and a fully specified multi-sector Ricardian model, we estimate productivities at the sector level and examine how they evolve over time in both developed and developing countries. We find that in both country groups, comparative advantage has become weaker: productivity grew systematically faster in sectors that were initially at the greater comparative disadvantage. The global welfare implications of this phenomenon are significant. Relative to the counterfactual scenario in which an individual country's comparative advantage remained the same as in the 1960s, and technology in all sectors grew at the same country-specific average rate, welfare today is 1.9% lower for the median country. The welfare impact varies greatly across countries, ranging from −0.5% to +6% among OECD countries, and from −9% to +27% among non-OECD countries. Contrary to frequently expressed concerns, changes in developing countries' comparative advantage had virtually no impact on welfare in the developed countries.

JEL Classifications: F11, F43, O33, O47

Keywords: technological change, sectoral TFP, Ricardian models of trade, welfare

*We are grateful to Costas Arkolakis, Alan Deardorff, Chris House, Francesc Ortega, Dmitriy Stolyarov, Linda Tesar, Kei-Mu Yi, and seminar participants at the University of Michigan, University of Toronto, NYU, the World Bank, the 2010 NBER IFM Fall meetings, 2011 ESSIM (Gerzensee), and 2011 SED (Gent) for helpful suggestions, and to Andrew McCallum and Lin Ma for excellent research assistance. E-mail (URL): alev@umich.edu (http://alevchenko.com), jzhang@umich.edu (http://www-personal.umich.edu/~jzhang/).

1 Introduction

How does technology evolve over time? This question is important in a variety of contexts, most notably in economic growth and international trade. Much of the economic growth literature focuses on *absolute* technological differences between countries. In the context of the one-sector model common in this literature, technological progress is unambiguously beneficial. Indeed, one reading of the growth literature is that most of the cross-country income differences are accounted for by technology, broadly construed (Klenow and Rodríguez-Clare 1997, Hall and Jones 1999).

By contrast, the Ricardian tradition in international trade emphasizes *relative* technological differences as the reason for international exchange and gains from trade. In the presence of multiple industries and comparative advantage, the welfare consequences of technological improvements depend crucially on which sectors experience productivity growth. For instance, it is well known that when productivity growth is biased towards sectors in which a country has a comparative disadvantage, the country and its trading partners may experience a welfare loss, relative to the alternative under which growth is balanced across sectors. Plainly, greater *relative* technology differences lead to larger gains from trade, and thus welfare is reduced when countries become more similar to each other. This result goes back to at least Hicks (1953), and has been reiterated recently by Samuelson (2004) in the context of productivity growth in developing countries.[1]

To fully account for the impact of technological progress on economic outcomes, we must understand not only the changes in average country-level TFP, but also how relative technology evolves across sectors. Or, in the vocabulary of international trade, it is important to know what happens to both absolute and comparative advantage. However, until now the literature has focused almost exclusively on estimating differences in technology at the country level. This paper examines the evolution of comparative advantage over time and its welfare consequences. Using a large-scale industry-level dataset on production and bilateral trade, spanning 75 countries, 19 manufacturing sectors, and 5 decades, we estimate productivity in each country, sector, and decade, and document the changes in comparative advantage between the 1960s and today. We then use these estimates together with a multi-sector Ricardian model of production and trade to quantify the welfare consequences of the patterns seen in the data.

Our main results can be summarized as follows. First, we find strong evidence that comparative advantage has become weaker. Controlling for the average productivity growth of all sectors in a country, sectors that were at the greater initial comparative disadvantage grew systematically faster. This effect is present in all time periods, and is similar in magnitude in both developed and developing countries. The speed of convergence in sectoral productivities implied by the estimates

[1]Other papers that explore technological change in Ricardian models are, among many others, Jones (1979), Krugman (1979), Brezis, Krugman and Tsiddon (1993), and Hymans and Stafford (1995).

is about 25% per decade.[2]

Second, the welfare impact of changes in comparative advantage is large. Our quantitative exercise begins by solving the full model under the actually observed pattern of comparative advantage, and computing welfare in each country in the 2000s under this baseline case. We then compare the baseline welfare to a counterfactual scenario in which an individual country's sectoral productivities grow at the same average rate between the 1960s and the 2000s, but its comparative advantage remains as it was in the 1960s. Because we allow average productivity to grow, this exercise isolates the role of changes in comparative – as opposed to absolute – advantage in welfare. The counterfactual also keeps the rest of the world's productivities same as in the data, and thus reveals the impact of changes in a single country's comparative advantage on its own welfare.

For the median country, welfare today is 1.9% lower than it would have been had comparative advantage remained unchanged since the 1960s. This median welfare impact corresponds to roughly 40% of the median gains from trade relative to autarky, 4.5%, implied by the model. Moreover, there is a great deal of variation around this average: the percentage difference between welfare under this counterfactual and the baseline ranges from -0.5% to $+6\%$ among OECD countries, and from -9% to $+27\%$ among non-OECD countries. Lower average welfare is exactly what theory would predict, given the empirical result that a typical country's comparative advantage has become weaker over this period. Indeed, we find that countries with a more pronounced weakening of comparative advantage tend to experience larger welfare losses, and countries whose comparative advantage strengthened tend to gain in welfare.

We next examine how each country's welfare is affected by technological change in its trading partners. It is sometimes suggested that changes in comparative advantage in developing countries can reduce welfare in developed ones (see Samuelson 2004, for a recent example). We evaluate this concern quantitatively, with the following two counterfactual exercises. In the first, we keep only the non-OECD comparative advantage fixed to the 1960s. In the second counterfactual, we keep only the comparative advantage in the OECD countries fixed as in the 1960s instead. Thus, the first (second) exercise reveals the global welfare changes that are due to the changes in comparative advantage in the non-OECD (OECD) only.

The main finding is that observed changes in developing countries' comparative advantage have virtually no impact on the OECD: the median welfare impact is zero, with a very narrow range of -0.2% to $+0.6\%$. This result is due in part to the fact that trade between the OECD countries

[2] Though a comprehensive investigation of theoretical mechanisms behind this finding is beyond the scope of this paper, we can conjecture that these empirical results are not supportive of learning-by-doing theories of comparative advantage (Krugman 1987, Young 1991), since these theories typically imply that productivity rises fastest in export sectors, and thus comparative advantage becomes more pronounced over time. Rather, our results are consistent with theories of "defensive innovation" in response to import competition (Bloom, Romer and Van Reenen 2010, Bloom, Draca and Van Reenen 2011).

still accounts for the majority of world trade, and thus the OECD countries are almost always each others' largest trading partners. Welfare in the non-OECD is also not affected by changes in the OECD comparative advantage, but due to a different mechanism. Closer inspection of the productivity estimates reveals, not surprisingly, that the frontier productivity in all sectors and all decades belongs to the OECD countries. Thus, from the perspective of a non-OECD country trading with the OECD as a group, it essentially always faces trade with the world frontier. While across decades, the particular OECD countries that occupy the frontier position may change, for a typical developing country these changes in comparative advantage in individual OECD countries turn out not to matter much.

To estimate productivity, the paper extends the methodology developed by Eaton and Kortum (2002) to a multi-sector framework. It is important to emphasize the advantages of our approach relative to the standard neoclassical methodology of computing measured TFP. The basic difficulty in directly measuring sectoral TFP in a large sample of countries and over time is the lack of comparable data on real sectoral output and inputs.[3] By contrast, our procedure uses information on bilateral trade, and thus dramatically expands the set of countries, sectors, and time periods for which productivity can be estimated. We follow the insight of Eaton and Kortum (2002) that trade flows contain information on productivity. Intuitively, if controlling for the typical gravity determinants of trade, a country spends relatively more on domestically produced goods in a particular sector, it is revealed to have either a high relative productivity or a low relative unit cost in that sector. We then use data on factor and intermediate input prices to net out the role of factor costs, yielding an estimate of relative productivity.

In addition, our approach extends the basic multi-sector Eaton-Kortum framework to incorporate many features that are important for reliably estimating underlying technology: multiple factors of production (labor and capital), differences in factor and intermediate input intensities across sectors, a realistic input-output matrix between the sectors, both inter- and intra-sectoral trade, and a non-traded sector. Finally, because our approach allows for international trade driven by both Ricardian and Heckscher-Ohlin forces, it takes explicit account of each country's participation in exports and imports, both of the final output, and of intermediate inputs used in production.

We are not the first to use international trade data to estimate technology parameters. Eaton and Kortum (2002) and Waugh (2010) perform this analysis in a one-sector model at a point in time, an exercise informative of the cross-section of countries' overall TFP but not their com-

[3]To our knowledge, the most comprehensive database that can be used to measure sectoral TFP on a consistent basis across countries and time is the OECD Structural Analysis (STAN) database. It contains the required information on only 12 developed countries for the period 1970-2008 in the best of cases, but upon closer inspection it turns out that the time and sectoral coverage is poor even in that small set of countries. Appendix A builds measured TFPs using the STAN database, and compares them to our estimates. There is a high positive correlation between the two, providing additional support for the validity of the estimates in this paper.

parative advantage.[4] Shikher (2004, 2005, 2011) and Costinot, Donaldson and Komunjer (2011) estimate sectoral technology for OECD countries, while Caliendo and Parro (2010) analyze the impact of NAFTA in a multi-sector Eaton-Kortum model. Hsieh and Ossa (2011) examine the global welfare impact of sector-level productivity growth in China between 1993 and 2005, focusing on the uneven growth across sectors. Chor (2010) relates Ricardian productivity differences to observable characteristics of countries, such as institutions and financial development. Relative to existing contributions, we extend the multi-sector approach to a much greater set of countries, and, most importantly, over time. This allows us, for the first time, to examine not only the global cross-section of productivities, but also its evolution over the past 5 decades and the welfare implications of those changes.

Changes in productivity at the sector level have received comparatively less attention in the literature. Bernard and Jones (1996a, 1996b) use production data to study convergence of measured TFP in a sample of 15 OECD countries and 8 sectors. Proudman and Redding (2000) and Hausmann and Klinger (2007) examine changes in countries' revealed comparative advantage and how these are related to initial export patterns. Our paper is the first to use a fully specified model of production and trade to estimate changes in technology. In addition, we greatly expand the sample of countries and years relative to these studies.

Finally, our paper is related to the literature that documents the time evolution of diversification indices, be it of production (e.g. Imbs and Wacziarg 2003), or trade (e.g. Carrère, Cadot and Strauss-Kahn 2009). These studies typically find that countries have a tendency to diversify their production and exports as they grow, at least until they become quite developed. Our findings of weakening comparative advantage are consistent with greater diversification. Unlike diversification indices, which have no structural interpretation, our approach makes this phenomenon more precise, by calculating the magnitudes of technology changes responsible for the observed changes in diversification.[5]

The rest of the paper is organized as follows. Section 2 lays out the theoretical framework. Section 3 presents the estimation procedure and the data. Section 4 describes the patterns of the evolution of comparative advantage over time, and presents the main econometric results of the paper on relative convergence. Section 5 examines the welfare implications of the observed evolution of comparative advantage. Section 6 concludes.

[4] Finicelli, Pagano and Sbracia (2009b) estimate the evolution of overall manufacturing TFP between 1985 and 2002 using a one-sector Eaton and Kortum model.

[5] Our paper is also related to the literature on international technology diffusion, surveyed by Keller (2004). While we document large and systematic changes in technology over time, our approach is, for now, silent on the mechanisms behind these changes.

2 Theoretical Framework

The world is comprised of N countries, indexed by n and i, and $J+1$ sectors, indexed by j and k. There are two factors of production, labor (L) and capital (K). Each sector produces a continuum of goods. The first J sectors are tradeable subject to barriers to trade, and the $J+1$-th sector is nontradeable. Both capital and labor are mobile across sectors and immobile across countries. Trade is balanced each period. We suppress the time index for the ease of notation.

2.1 The Environment

Period utility of the representative consumer in country n is homothetic, given by

$$U_n = \left(\sum_{j=1}^{J} \omega_j^{\frac{1}{\eta}} \left(Y_n^j \right)^{\frac{\eta-1}{\eta}} \right)^{\frac{\eta}{\eta-1}\xi_n} \left(Y_n^{J+1} \right)^{1-\xi_n}, \tag{1}$$

where ξ_n denotes the Cobb-Douglas weight for the tradeable sector composite good, η is the elasticity of substitution between the tradeable sectors, ω_j is the taste parameter for tradeable sector j, Y_n^{J+1} is the nontradeable-sector composite good, and Y_n^j is the composite good in tradeable sector j.[6] The budget constraint (or the resource constraint) of the consumer is given by

$$\sum_{j=1}^{J+1} p_n^j Y_n^j = w_n L_n + r_n K_n, \tag{2}$$

where p_n^j is the price of the sector j composite, K_n and L_n are the exogenous capital and labor endowments, and w_n and r_n are the wage rate and the rental return of capital, respectively. The consumption price index in country n is thus:

$$P_n = B_n \left(\sum_{j=1}^{J} \omega_j (p_n^j)^{1-\eta} \right)^{\frac{1}{1-\eta}\xi_n} (p_n^{J+1})^{1-\xi_n},$$

where $B_n = \xi_n^{-\xi_n}(1-\xi_n)^{-(1-\xi_n)}$.

Output in each sector j is produced using a CES production function that aggregates a continuum of varieties $q \in [0,1]$ unique to each sector:

$$Q_n^j = \left[\int_0^1 Q_n^j(q)^{\frac{\varepsilon-1}{\varepsilon}} dq \right]^{\frac{\varepsilon}{\varepsilon-1}},$$

[6] In the quantitative exercise that follows, the share of non-tradeables ξ_n will vary over time as well as across countries, to capture the positive relationship between income and the non-tradeables consumption share observed in the data.

where ε denotes the elasticity of substitution across varieties q, Q_n^j is the total output of sector j in country n, and $Q_n^j(q)$ is the amount of variety q that is used in production in sector j and country n. It is well known that the price of sector j's output is given by:

$$p_n^j = \left[\int_0^1 p_n^j(q)^{1-\varepsilon} dq\right]^{\frac{1}{1-\varepsilon}}.$$

Producing one unit of good q in sector j in country n requires $\frac{1}{z_n^j(q)}$ input bundles. The cost of an input bundle is:

$$c_n^j = \left(w_n^{\alpha_j} r_n^{1-\alpha_j}\right)^{\beta_j} \left(\prod_{k=1}^{J+1} \left(p_n^k\right)^{\gamma_{k,j}}\right)^{1-\beta_j}.$$

That is, production in sector j requires labor, capital, and a bundle of intermediate inputs, coming from all sectors $k = 1, ..., J+1$. The value-added based labor intensity is given by α_j, while the share of value added in total output is given by β_j. Both of these vary by sector. The weights on inputs from other sectors, $\gamma_{k,j}$ vary by output industry j as well as input industry k.

Productivity $z_n^j(q)$ for each $q \in [0,1]$ in each sector j is equally available to all agents in country n, and product and factor markets are perfectly competitive. Following Eaton and Kortum (2002, henceforth EK), the productivity draw $z_n^j(q)$ is random and comes from the Fréchet distribution that has the cumulative distribution function

$$F_n^j(z) = e^{-T_n^j z^{-\theta}}.$$

In this distribution, the absolute advantage term T_n^j varies by both country and sector, with higher values of T_n^j implying higher average productivity draws in sector j in country n. The parameter θ captures dispersion, with larger values of θ implying smaller dispersion in draws.

The cost of producing one unit of good q in sector j and country n is $c_n^j/z_n^j(q)$. International trade is subject to iceberg costs: in order for one unit of good q produced in sector j to arrive at country n from country i, $d_{ni}^j > 1$ units of the good must be shipped. We normalize $d_{nn}^j = 1$ for each country n in each tradeable sector j. Note that the trade costs will vary by destination pair, by sector, and by time, and in general will not be symmetric: d_{ni}^j need not equal d_{in}^j. Under perfect competition, the price at which country i can supply tradeable good q in sector j to country n is equal to:

$$p_{ni}^j(q) = \left(\frac{c_i^j}{z_i^j(q)}\right) d_{ni}^j.$$

Buyers of each good q in tradeable sector j in country n will select to buy from the cheapest source country. Thus, the price actually paid for this good in country n will be:

$$p_n^j(q) = \min_{i=1,\dots,N} \left\{ p_{ni}^j(q) \right\}.$$

Following the standard EK approach, define the "multilateral resistance" term

$$\Phi_n^j = \sum_{i=1}^{N} T_i^j \left(c_i^j d_{ni}^j \right)^{-\theta}.$$

This value summarizes, for country n, the access to production technologies in sector j. Its value will be higher if in sector j, country n's trading partners have high productivity (T_i^j) or low cost (c_i^j). It will also be higher if the trade costs that country n faces in this sector are low. Standard steps lead to the familiar result that the probability of importing good q from country i, π_{ni}^j is equal to the share of total spending on goods coming from country i, X_{ni}^j / X_n^j, and is given by:

$$\frac{X_{ni}^j}{X_n^j} = \pi_{ni}^j = \frac{T_i^j \left(c_i^j d_{ni}^j \right)^{-\theta}}{\Phi_n^j}.$$

In addition, the price of good j in country n is simply

$$p_n^j = \Gamma \left(\Phi_n^j \right)^{-\frac{1}{\theta}}, \tag{3}$$

where $\Gamma = \left[\mathbf{\Gamma} \left(\frac{\theta + 1 - \varepsilon}{\theta} \right) \right]^{\frac{1}{1-\varepsilon}}$, with $\mathbf{\Gamma}$ the Gamma function.

2.2 Equilibrium

The **competitive equilibrium** of this model world economy consists of a set of prices, allocation rules, and trade shares such that (i) given the prices, all firms' inputs satisfy the first-order conditions, and their output is given by the production function; (ii) given the prices, the consumers' demand satisfies the first-order conditions; (iii) the prices ensure the market clearing conditions for labor, capital, tradeable goods and nontradeable goods; (iv) trade shares ensure balanced trade for each country.

The set of prices includes the wage rate w_n, the rental rate r_n, the sectoral prices $\{p_n^j\}_{j=1}^{J+1}$, and the aggregate price P_n in each country n. The allocation rules include the capital and labor allocation across sectors $\{K_n^j, L_n^j\}_{j=1}^{J+1}$, final consumption demand $\{Y_n^j\}_{j=1}^{J+1}$, and total demand $\{Q_n^j\}_{j=1}^{J+1}$ (both final and intermediate goods) for each sector. The trade shares include the expenditure share π_{ni}^j in country n on goods coming from country i in sector j.

Characterization of Equilibrium

Given the set of prices $\{w_n, r_n, P_n, \{p_n^j\}_{j=1}^{J+1}\}_{n=1}^N$, we first characterize the optimal allocations from final demand. Consumers maximize utility (1) subject to the budget constraint (2). The first order conditions associated with this optimization problem imply the following final demand:

$$p_n^j Y_n^j = \xi_n (w_n L_n + r_n K_n) \frac{\omega_j (p_n^j)^{1-\eta}}{\sum_{k=1}^J \omega_k (p_n^k)^{1-\eta}}, \text{ for all } j = \{1, .., J\} \tag{4}$$

and

$$p_n^{J+1} Y_n^{J+1} = (1 - \xi_n)(w_n L_n + r_n K_n).$$

We next characterize the production and factor allocations across the world. Let Q_n^j denote the total sectoral demand in country n and sector j. Q_n^j is used for both final consumption and intermediate inputs in domestic production of all sectors. That is,

$$p_n^j Q_n^j = p_n^j Y_n^j + \sum_{k=1}^J (1 - \beta_k) \gamma_{j,k} \left(\sum_{i=1}^N \pi_{in}^k p_i^k Q_i^k \right) + (1 - \beta_{J+1}) \gamma_{j,J+1} p_n^{J+1} Q_n^{J+1}$$

for tradeable sectors $j = 1, ..., J$, and

$$p_n^{J+1} Q_n^{J+1} = p_n^{J+1} Y_n^{J+1} + \sum_{k=1}^{J+1} (1 - \beta_k) \gamma_{j,k} p_n^k Q_n^k$$

in the nontradeable sector. That is, total expenditure in sector $j = 1, ..., J$ of country n, $p_n^j Q_n^j$, is the sum of (i) domestic final consumption expenditure $p_n^j Y_n^j$; (ii) expenditure on sector j goods as intermediate inputs in all the traded sectors $\sum_{k=1}^J (1-\beta_k) \gamma_{j,k} (\sum_{i=1}^N \pi_{in}^k p_i^k Q_i^k)$, and (iii) expenditure on the j's sector intermediate inputs in the domestic non-traded sector $(1-\beta_{J+1})\gamma_{j,J+1} p_n^{J+1} Q_n^{J+1}$. These market clearing conditions summarize the two important features of the world economy captured by our model: complex international production linkages, as much of world trade is in intermediate inputs, and a good crosses borders multiple times before being consumed (Hummels, Ishii and Yi 2001); and two-way input linkages between the tradeable and the nontradeable sectors.

In each tradeable sector j, some goods q are imported from abroad and some goods q are exported to the rest of the world. Country n's exports in sector j are given by $EX_n^j = \sum_{i=1}^N \mathbb{I}_{i \neq n} \pi_{in}^j p_i^j Q_i^j$, and its imports in sector j are given by $IM_n^j = \sum_{i=1}^N \mathbb{I}_{i \neq n} \pi_{ni}^j p_n^j Q_n^j$, where $\mathbb{I}_{i \neq n}$ is the indicator function. The total exports of country n are then $EX_n = \sum_{j=1}^J EX_n^j$, and total imports are $IM_n = \sum_{j=1}^J IM_n^j$. Trade balance requires that for any country n, $EX_n - IM_n = 0$.

We now study the factor allocations across sectors. The total production revenue in tradeable sector j in country n is given by $\sum_{i=1}^N \pi_{in}^j p_i^j Q_i^j$. The optimal sectoral factor allocations in country

n and tradeable sector j must thus satisfy

$$\sum_{i=1}^{N} \pi_{in}^{j} p_i^j Q_i^j = \frac{w_n L_n^j}{\alpha_j \beta_j} = \frac{r_n K_n^j}{(1-\alpha_j)\beta_j}.$$

For the nontradeable sector $J+1$, the optimal factor allocations in country n are simply given by

$$p_n^{J+1} Q_n^{J+1} = \frac{w_n L_n^{J+1}}{\alpha_{J+1}\beta_{J+1}} = \frac{r_n K_n^{J+1}}{(1-\alpha_{J+1})\beta_{J+1}}.$$

Finally, the feasibility conditions for factors are given by, for any n,

$$\sum_{j=1}^{J+1} L_n^j = L_n \text{ and } \sum_{j=1}^{J+1} K_n^j = K_n.$$

Given all of the model parameters, factor endowments, trade costs, and productivities, the model is solved using the algorithm described in Appendix B.

3 Estimating Model Parameters

This section estimates the sector-level technology parameters T_n^j for a large set of countries and 5 decades, in three steps. First, we estimate the technology parameters in the tradeable sectors relative to the U.S. using data on sectoral output and bilateral trade. The procedure relies on fitting a structural gravity equation implied by the model. Intuitively, if controlling for the typical gravity determinants of trade, a country spends relatively more on domestically produced goods in a particular sector, it is revealed to have either a high relative productivity or a low relative unit cost in that sector. We then use data on factor and intermediate input prices to net out the role of factor costs, yielding an estimate of relative productivity. This step also produces estimates of bilateral trade costs at the sectoral level over time. Second, we estimate the technology parameters in the tradeable sectors for the U.S.. This procedure requires directly measuring TFP at the sectoral level using data on real output and inputs, and then correcting measured TFP for selection due to trade. The taste parameters for all tradeable sectors ω_j are also calibrated in this step. Third, we calibrate the nontradeable technology for all countries using the first-order condition of the model and the relative prices observed in the data.

The calibration of the remaining parameters is more straightforward. Some parameters – $\alpha_j, \beta_j, \gamma_{k,j}, \xi_n, L_n$, and K_n – come directly from the data.[7] For a small number of parameters – θ, η, and ε – we take values estimated elsewhere in the literature. We discuss the data sources used to calibrate these parameters in detail in Section 3.3.

[7] While α_j, β_j, and $\gamma_{k,j}$ are time-invariant, ξ_n, L_n, and K_n will vary within country over time.

3.1 Tradeable Sector Relative Technology

We now focus on the tradeable sectors. Following the standard EK approach, first divide trade shares by their domestic counterpart:

$$\frac{\pi^j_{ni}}{\pi^j_{nn}} = \frac{X^j_{ni}}{X^j_{nn}} = \frac{T^j_i \left(c^j_i d^j_{ni}\right)^{-\theta}}{T^j_n \left(c^j_n\right)^{-\theta}},$$

which in logs becomes:

$$\ln\left(\frac{X^j_{ni}}{X^j_{nn}}\right) = \ln\left(T^j_i (c^j_i)^{-\theta}\right) - \ln\left(T^j_n (c^j_n)^{-\theta}\right) - \theta \ln d^j_{ni}.$$

Let the (log) iceberg costs be given by the following expression:

$$\ln d^j_{ni} = d^j_k + b^j_{ni} + CU^j_{ni} + RTA^j_{ni} + ex^j_i + \nu^j_{ni},$$

where d^j_k is an indicator variable for a distance interval. Following EK, we set the distance intervals, in miles, to [0, 350], [350, 750], [750, 1500], [1500, 3000], [3000, 6000], [6000, maximum). Additional variables are whether the two countries share a common border (b^j_{ni}), belong to a currency union (CU^j_{ni}), or to a regional trade agreement (RTA^j_{ni}). Following the arguments in Waugh (2010), we include an exporter fixed effect ex^j_i. Finally, there is an error term ν^j_{ni}. Note that all the variables have a sector superscript j: we allow all the trade cost proxy variables to affect true iceberg trade costs d^j_{ni} differentially across sectors. There is a range of evidence that trade volumes at sector level vary in their sensitivity to distance or common border (see, among many others, Do and Levchenko 2007, Berthelon and Freund 2008).

This leads to the following final estimating equation:

$$\ln\left(\frac{X^j_{ni}}{X^j_{nn}}\right) = \underbrace{\ln\left(T^j_i (c^j_i)^{-\theta}\right) - \theta ex^j_i}_{\text{Exporter Fixed Effect}} \underbrace{- \ln\left(T^j_n \left(c^j_n\right)^{-\theta}\right)}_{\text{Importer Fixed Effect}} \underbrace{-\theta d^j_k - \theta b^j_{ni} - \theta CU^j_{ni} - \theta RTA^j_{ni}}_{\text{Bilateral Observables}} \underbrace{-\theta \nu^j_{ni}}_{\text{Error Term}}.$$

Estimating this relationship will thus yield, for each country, an estimate of its technology-cum-unit-cost term in each sector j, $T^j_n (c^j_n)^{-\theta}$, which is obtained by exponentiating the importer fixed effect.[8] The available degrees of freedom imply that these estimates are of each country's

[8] A standard feature of this procedure is that the trade shares are logged, so that the zero bilateral import flows are dropped from the estimation sample (Eaton and Kortum 2002, Waugh 2010). Unfortunately, our large-scale model cannot be tractably enriched to explicitly account for zeros in trade and at the same time retain the structural interpretation linking the fixed effects to underlying productivity. However, we can check the ex-post performance of

$T_n^j (c_n^j)^{-\theta}$ relative to a reference country, which in our estimation is the United States. We denote this estimated value by S_n^j:

$$S_n^j = \frac{T_n^j}{T_{us}^j} \left(\frac{c_n^j}{c_{us}^j} \right)^{-\theta},$$

where the subscript us denotes the United States. It is immediate from this expression that estimation delivers a convolution of technology parameters T_n^j and cost parameters c_n^j. Both will of course affect trade volumes, but we would like to extract technology T_n^j from these estimates. In order to do that, we follow the approach of Shikher (2004). In particular, for each country n, the share of total spending going to home-produced goods is given by

$$\frac{X_{nn}^j}{X_n^j} = T_n^j \left(\frac{\Gamma c_n^j}{p_n^j} \right)^{-\theta}.$$

Dividing by its U.S. counterpart yields:

$$\frac{X_{nn}^j / X_n^j}{X_{us,us}^j / X_{us}^j} = \frac{T_n^j}{T_{us}^j} \left(\frac{c_n^j}{c_{us}^j} \frac{p_{us}^j}{p_n^j} \right)^{-\theta} = S_n^j \left(\frac{p_{us}^j}{p_n^j} \right)^{-\theta},$$

and thus the ratio of price levels in sector j relative to the U.S. becomes:

$$\frac{p_n^j}{p_{us}^j} = \left(\frac{X_{nn}^j / X_n^j}{X_{us,us}^j / X_{us}^j} \frac{1}{S_n^j} \right)^{\frac{1}{\theta}}. \tag{5}$$

The entire right-hand side of this expression is either observable or estimated. Thus, we can impute the price levels relative to the U.S. in each country and each tradeable sector.

The cost of the input bundles relative to the U.S. can be written as:

$$\frac{c_n^j}{c_{us}^j} = \left(\frac{w_n}{w_{us}} \right)^{\alpha_j \beta_j} \left(\frac{r_n}{r_{us}} \right)^{(1-\alpha_j)\beta_j} \left(\prod_{k=1}^{J} \left(\frac{p_n^k}{p_{us}^k} \right)^{\gamma_{k,j}} \right)^{1-\beta_j} \left(\frac{p_n^{J+1}}{p_{us}^{J+1}} \right)^{\gamma_{J+1,j}(1-\beta_j)}.$$

Using information on relative wages, returns to capital, price in each tradeable sector from (5),

the estimated model with respect to zeros by solving the full model, and computing within the model the sum of the π_{ni}^j's in the importer-exporter-sector observations that are zeros in the actual data. We can then examine whether these observations account for large shares of absorption inside the model. If the resulting numbers are large, then the quantitative model predicts substantial trade flows where in reality there is zero. However, if these numbers are small, the model predicts very small flows where the actual flows are zero, providing a good approximation to the data even though productivities are estimated dropping zero trade. The results of this exercise are reported in Appendix Table A3. Observations for which actual data exhibit zero trade account for a tiny share of overall absorption in our quantitative model: in each decade, these observations add up to on average 0.1% of the total absorption. Breaking down across sectors and decades, we see that in none of the individual sectors or decades, this share is greater than 0.8%, and for the vast majority sector×decade observations, the numbers are at or below 0.2%. We conclude from this exercise that in spite of ignoring the zero trade observations in estimation, our quantitative model is quite close to the data when it comes to small/zero trade flows.

and the nontradeable sector price relative to the U.S., we can thus impute the costs of the input bundles relative to the U.S. in each country and each sector. Armed with those values, it is straightforward to back out the relative technology parameters:

$$\frac{T_n^j}{T_{us}^j} = S_n^j \left(\frac{c_n^j}{c_{us}^j} \right)^{\theta}.$$

3.2 Complete Estimation

So far we have estimated the levels of technology of the tradeable sectors relative to the United States. To complete our estimation, we still need to find (i) the levels of T for the tradeable sectors in the United States; (ii) the taste parameters ω_j, and (iii) the nontradeable technology levels for all countries.

To obtain (i), we use the NBER-CES Manufacturing Industry Database for the U.S. (Bartelsman and Gray 1996). We start by measuring the observed TFP levels for the tradeable sectors in the U.S.. The form of the production function gives

$$\ln Z_{us}^j = \ln \Lambda_{us}^j + \beta_j \alpha_j \ln L_{us}^j + \beta_j (1-\alpha_j) \ln K_{us}^j + (1-\beta_j) \sum_{k=1}^{J+1} \gamma_{k,j} \ln M_{us}^{k,j}, \tag{6}$$

where Λ^j denotes the measured TFP in sector j, Z^j denotes the output, L^j denotes the labor input, K^j denotes the capital input, and $M^{k,j}$ denotes the intermediate input from sector k. The NBER-CES Manufacturing Industry Database offers information on output, and inputs of labor, capital, and intermediates, along with deflators for each. Thus, we can estimate the observed TFP level for each manufacturing tradeable sector using the above equation.

If the United States were a closed economy, the observed TFP level for sector j would be given by $\Lambda_{us}^j = (T_{us}^j)^{\frac{1}{\theta}}$. In the open economies, the goods with inefficient domestic productivity draws will not be produced and will be imported instead. Thus, international trade and competition introduce selection in the observed TFP level, as demonstrated by Finicelli, Pagano and Sbracia (2009a). We thus use the model to back out the true level of T_{us}^j of each tradeable sector in the United States. Here we follow Finicelli et al. (2009a) and use the following relationship:

$$(\Lambda_{us}^j)^{\theta} = T_{us}^j + \sum_{i \neq us} T_i^j \left(\frac{c_i^j d_{us,i}^j}{c_{us}^j} \right)^{-\theta}.$$

Thus, we have

$$(\Lambda_{us}^j)^{\theta} = T_{us}^j \left[1 + \sum_{i \neq us} \frac{T_i^j}{T_{us}^j} \left(\frac{c_i^j d_{us,i}^j}{c_{us}^j} \right)^{-\theta} \right] = T_{us}^j \left[1 + \sum_{i \neq us} S_i^j \left(d_{us,i}^j \right)^{-\theta} \right]. \tag{7}$$

This equation can be solved for underlying technology parameters T_{us}^j in the U.S., given estimated observed TFP Λ_{us}^j, and all the S_i^j's and $d_{us,i}^j$'s estimated in the previous subsection.

To estimate the taste parameters $\{\omega_j\}_{j=1}^J$, we use information on final consumption shares in the tradeable sectors in the U.S.. We start with a guess of $\{\omega_j\}_{j=1}^J$ and find sectoral prices p_n^k as follows. For an initial guess of sectoral prices, we compute the tradeable sector aggregate price and the nontradeable sector price using the data on the relative prices of nontradeables to tradeables. Using these prices, we calculate sectoral unit costs and Φ_n^j's, and update prices according to equation (3), iterating until the prices converge. We then update the taste parameters according to equation (4), using the data on final sectoral expenditure shares in the U.S.. We normalize the vector of ω_j's to have a sum of one, and repeat the above procedure until the values for the taste parameters converge.

Finally, we estimate the nontradeable sector TFP using the relative prices. In the model, the nontradeable sector price is given by

$$p_n^{J+1} = \Gamma(T_n^{J+1})^{-\frac{1}{\theta}} c_n^{J+1}.$$

Since we know the aggregate price level in the tradeable sector p_n^T, c_n^{J+1}, and the relative price of nontradeables (which we take from the data), we can back out T_n^{J+1} from the equation above for all countries.

3.3 Data Description and Implementation

In order to carry out estimation, we assemble data on production and trade for a sample of up to 75 countries, 19 manufacturing sectors, and spanning 5 decades, from the 1960s to the 2000s. Production data come from the 2009 UNIDO Industrial Statistics Database, which reports output, value added, employment, and wage bills at roughly 2-digit ISIC Revision 3 level of disaggregation for the period 1962-2007 in the best of cases. The corresponding trade data comes from the COMTRADE database compiled by the United Nations. The trade data are collected at the 4-digit SITC level, and aggregated up to the 2-digit ISIC level using a concordance developed by the authors. Production and trade data were extensively checked for quality, and a number of countries were discarded due to poor data quality. In addition, in less than 5% of country-year-sector observations, the reported total output was below total exports, and thus had to be imputed based on earlier values and the evolution of exports. Appendix Table A1 lists the countries used in the analysis along with the time periods for which data are available for each country, and Appendix Table A2 lists the sectors along with the key parameter values for each sector: α_j, β_j, the share of nontradeable inputs in total inputs $\gamma_{J+1,j}$, and the taste parameter ω_j.

The distance and common border variables are obtained from the comprehensive geography database compiled by CEPII. Information on regional trade agreements comes from the RTA database maintained by the WTO. The currency union indicator comes from Rose (2004), and was updated for the post-2000 period using publicly available information (such as the membership in the Euro area, and the dollarization of Ecuador and El Salvador).

In addition to providing data on output for gravity estimation, the UNIDO data are used to estimate production function parameters α_j and β_j. To compute α_j for each sector, we calculate the share of the total wage bill in value added, and take a simple median across countries (taking the mean yields essentially the same results). To compute β_j, take the median of value added divided by total output.

The intermediate input coefficients $\gamma_{k,j}$ are obtained from the Direct Requirements Table for the United States. We use the 1997 Benchmark Detailed Make and Use Tables (covering approximately 500 distinct sectors), as well as a concordance to the ISIC Revision 3 classification to build a Direct Requirements Table at the 2-digit ISIC level. The Direct Requirements Table gives the value of the intermediate input in row k required to produce one dollar of final output in column j. Thus, it is the direct counterpart of the input coefficients $\gamma_{k,j}$. Note that we assume these to be the same in all countries.[9] In addition, we use the U.S. I-O matrix to obtain the shares of total final consumption expenditure going to each sector, which we use to pin down taste parameters ω_j in traded sectors $1, ..., J$; as well as α_{J+1} and β_{J+1} in the nontradeable sector, which cannot be obtained from UNIDO.[10]

The computation of relative costs of the input bundle requires information on wages and the returns to capital. To compute wages, we divide the total manufacturing sector wage bill by total manufacturing employment in each country, and take that value relative to the U.S.. Consistent with the model, this procedure delivers wages that differ by country but not by sector.[11]

Obtaining information on the return to capital, r_n, is less straightforward, since it is not observable directly. In the baseline analysis, we assume that the wage-rental ratio is determined

[9]di Giovanni and Levchenko (2010) provide suggestive evidence that at such a coarse level of aggregation, Input-Output matrices are indeed similar across countries. We implemented two robustness checks so assess how restrictive is the assumption of a single I-O matrix across countries and over time. In the first, we collected U.S. I-O matrices for each decade going back to the 1960s, and re-calculated all the sectoral productivities using decade-specific matrices. For all the decades, the correlation between the productivities implied by the decade-specific matrices and the baseline values reported in the paper is extremely high: across all the countries in the sample, the median correlation is above 0.99, with the lowest correlation across all countries and decades of 0.91. Second, we collected country-specific I-O matrices from the GTAP database. Productivities computed based on country-specific I-O matrices were again very similar to the baseline values. In our sample of countries, the median correlation was 0.98, with all but 3 out of 75 countries having a correlation of 0.93 or above, and the minimum correlation of 0.65.

[10]The U.S. I-O matrix provides an alternative way of computing α_j and β_j. These parameters calculated based on the U.S. I-O table are very similar to those obtained from UNIDO, with the correlation coefficients between them above 0.85 in each case. The U.S. I-O table implies greater variability in α_j's and β_j's across sectors than does UNIDO.

[11]In less than 1% of country-decade observations, either the total wage bill or employment are missing from the UNIDO data. In those cases, the wage relative to the U.S. is proxied by the GDP per capita relative to the U.S.

by the aggregate capital-labor ratio through an aggregate market clearing condition: $r_n/w_n = ((1-\alpha)L_n)/(\alpha K_n)$, where α is the aggregate share of labor in GDP, which we set to 2/3.[12]

The price of nontradeables relative to the U.S., p_n^{J+1}/p_{us}^{J+1}, and the price of nontradeables relative to tradeables in each country, p_n^{J+1}/p_n^T, are computed using the detailed price data collected by the International Comparison of Prices Program (ICP). For a few countries and decades, these relative prices are extrapolated using a simple linear fit to log PPP-adjusted per capita GDP from the Penn World Tables 6.3 (Heston, Summers and Aten 2002).

The total labor force in each country, L_n, and the total capital stock, K_n, are obtained from the Penn World Tables 6.3. Following the standard approach in the literature (see, e.g. Hall and Jones 1999, Bernanke and Gürkaynak 2001, Caselli 2005), the total labor force is calculated from the data on the total GDP per capita and per worker.[13] The total capital is calculated using the perpetual inventory method that assumes a depreciation rate of 6%: $K_{n,t} = (1-0.06)K_{n,t-1} + I_{n,t}$, where $I_{n,t}$ is total investment in country n in period t. For most countries, investment data start in 1950, and the initial value of K_n is set equal to $I_{n,0}/(\gamma + 0.06)$, where γ is the average growth rate of investment in the first 10 years for which data are available.

In order to estimate the relative TFP's in the tradedable sectors in the U.S., we use the 2009 version of the NBER-CES Manufacturing Industry Database, that reports the total output, total input usage, employment, and capital stock, along with deflators for each of these in each sector. The data are available in the 6-digit NAICS classification for the period 1958 to 2005, and are converted into ISIC 2-digit sectors using a concordance developed by the authors. The procedure yields sectoral measured TFP's for the U.S. in each tradeable sector $j = 1, ..., J$ and each decade.

The share of expenditure on traded goods, ξ_n in each country and decade is sourced from Yi and Zhang (2010), who compile this information for 30 developed and developing countries. For countries unavailable in the Yi and Zhang data, values of ξ_n are imputed based on fitting a simple linear relationship to log PPP-adjusted per capita GDP from the Penn World Tables. In each decade, the fit of this simple linear relationship was typically quite good, with R^2's of 0.30 to 0.80 across decades.

Finally, for now we assume that the dispersion parameter θ does not vary across sectors. There are no reliable estimates of how it varies across sectors, and thus we do not model this variation. We pick the value of $\theta = 8.28$, which is the preferred estimate of EK.[14] It is important

[12] The return to capital will be affected by country characteristics other than capital abundance, such as the quality of the country's regulatory environment, corruption, and expropriation risk, among other factors. Indeed, Caselli and Feyrer (2007) document that the marginal product of capital is remarkably similar across a wide range countries. Alternatively, the return to capital will be the same in all countries under international capital mobility. None of the results below are affected if we assume instead that the return to capital, r_n, does not differ across countries.

[13] Using the variable name conventions in the Penn World Tables, $L_n = 1000 * pop * rgdpch/rgdpwok$.

[14] Shikher (2004, 2005, 2011), Burstein and Vogel (2009), and Eaton, Kortum, Neiman and Romalis (2010), among others, follow the same approach of assuming the same θ across sectors. Caliendo and Parro (2010) use tariff data

to assess how the results below are affected by the value of this parameter. One may be especially concerned about how the results change under lower values of θ. Lower θ implies greater within-sector heterogeneity in the random productivity draws. Thus, trade flows become less sensitive to the costs of the input bundles (c_i^j), and the gains from intra-sectoral trade become larger relative to the gains from inter-sectoral trade. We repeated the entire analysis in the paper assuming instead a value of $\theta = 4$, which has been advocated by Simonovska and Waugh (2010) and is at or near the bottom of the range that has been used in the literature. Overall, the results are remarkably similar. The correlation between estimated T_i^j's under $\theta = 4$ and the baseline is above 0.95, and there is actually somewhat greater variability in T_i^j's under $\theta = 4$. Appendix Tables A6 through A8 report the main econometric and quantitative results of the paper under this alternative value of θ. Comparing them to the baseline results, it is clear that the two are remarkably similar.

We choose the elasticity of substitution between broad sectors within the tradeable bundle, η, to be equal to 2. Since these are very large product categories, it is sensible that this elasticity would be relatively low. It is higher, however, than the elasticity of substitution between tradeable and nontradeable goods, which is set to 1 by the Cobb-Douglas assumption. The elasticity of substitution between varieties within each tradeable sector, ε, is set to 4.

All of the variables that vary over time are averaged for each decade, from the 1960s to the 2000s, and these decennial averages are used in the analysis throughout. Thus, our unit of time is a decade.

4 Evolution of Comparative Advantage

In this section, we describe the basic patterns in how estimated sector-level technology varies across countries and over time, focusing especially on whether comparative advantage has become stronger or weaker. Going through the steps described in Section 3.1 yields, for each country n, tradeable sector j, and decade, the state of technology relative to the U.S., T_n^j/T_{us}^j. Since the choice of the U.S. as the reference country is arbitrary, we present the stylized facts based not on each country's difference with respect to the U.S., but with respect to the global frontier. In each sector and decade, we select the 2 highest values of T_n^j/T_{us}^j, take their geometric mean, and label that the global frontier. We then re-normalize each country's technology parameter to be expressed relative to the frontier, rather than the U.S.. In addition, since mean productivity in

and triple differencing to estimate sector-level θ. However, their approach may impose too much structure and/or be dominated by measurement error: at times the values of θ they estimate are negative. In addition, in each sector the restriction that $\theta > \varepsilon - 1$ must be satisfied, and it is not clear whether Caliendo and Parro (2010)'s estimated sectoral θ's meet this restriction in every case. Our approach is thus conservative by being agnostic on this variation across sectors.

each sector is equal to $T^{1/\theta}$, we carry out the analysis on this value, rather than T.

4.1 Basic Patterns

Table 1 presents summary statistics for the OECD and non-OECD countries in each decade. The first column reports the mean distance to the frontier across all sectors in a country, a measure that can be thought of as *absolute advantage*. Not surprisingly, the OECD countries as a group catch up to the frontier between the 1960s and the 2000s, with productivities going up from 0.65 to 0.84 of the frontier value. The non-OECD countries' position shows no clear upward or downward pattern. The second column in each panel summarizes the magnitude of within-country differences in productivity across sectors. Namely, it reports the mean ratio of productivities in the two most productive sectors relative to the two least productive ones, by country group and decade. This measure can be thought of as *comparative advantage* across sectors. For the OECD, this measure is on the order of $1.4-1.5$, and decreasing monotonically over time. For the non-OECD countries, it fluctuates around 2, showing no clear trend. Not surprisingly, the non-OECD countries tend to have stronger comparative advantage.

The evolution of these averages over time masks a great deal of heterogeneity among countries. Table 2 reports top 10 and bottom 10 countries ranked according to how fast their average productivity changed relative to the frontier. The left panel presents the changes from the 1960s to 2000s, and the right panel from the 1980s to 2000s. Over the period 1960s-2000s, the countries that caught up to the frontier the fastest are for the most part peripheral OECD countries, such as Norway, Portugal, and Greece. Countries slowest to catch up (or fastest to fall behind) are developing countries, that surprisingly include two of the more successful East Asian economies, Thailand and Malaysia. This is of course not inconsistent with high rates of economic growth experienced by these countries. First, these are measures of average technology, and part of the growth in those countries would have been driven by factor accumulation. More importantly, these are measures of distance to the technological frontier. Thus, even if these countries experienced overall productivity growth, our procedure shows that the frontier grew even faster. Since the 1980s, the composition of countries changes somewhat, but the patterns are broadly similar.

In addition to absolute advantage, we can assess how the countries comparative advantage evolved. Table 3 reports the top 10 and bottom 10 countries in how much the dispersion in the country's technology across sectors changed. In particular, for each country and decade, we compute the coefficient of variation in $T^{1/\theta}$ across sectors, and record how much this coefficient of variation changed over time. Thus, larger negative changes imply greater reductions in productivity dispersion across sectors, and thus greater *relative* catch-up. Conversely, positive values imply that a country's comparative advantage has gotten stronger, as its productivity dispersion increased.

It is clear from comparing Tables 2 and 3 that absolute and relative convergence are closely related: most of the fastest converging countries on average are also those that catch up disproportionately in their weakest sectors. This can be due in part to the fact that the best sectors in those countries are already at the frontier, thus the only sectors that can catch up are the weak ones. However, the rankings are very similar if we instead do not normalize by the frontier, and assess the changes relative to a reference country. This way, there is no mechanical ceiling for a country's strongest sectors. Less obviously, the bottom countries tend to be similar as well. Thus, countries that fall behind the most on average also tend to experience greater dispersion across sectors: their weakest sectors fall disproportionately more than their strongest ones. Figure 1 presents the correlation between relative and absolute convergence graphically. There is a strong association between these two measures.

Table 4 reports the correlation coefficients between absolute and relative convergence measures, and the corresponding changes in real PPP-adjusted per capita income and overall trade openness, sourced from the Penn World Tables. In addition to the high positive correlation (0.61-0.64) between our two measures, the table reveals that neither is particularly strongly correlated with changes in income or openness. There is a positive correlation (around 0.25) between income growth and average convergence, while the correlation with relative convergence is close to zero and mildly negative. Growth in trade openness is actually negatively correlated with average convergence, and virtually uncorrelated with relative convergence. Figure 2 presents the scatterplots of absolute and relative convergence against income growth and openness.

Another important question is whether our estimates can be cross-validated using direct estimates of measured TFP. Appendix A estimates measured TFP using data on real output and inputs from the OECD Structural Analysis database. It is the most comprehensive database that contains the information required to estimate measured TFP on a consistent basis across countries and over time. Using both simple correlations and regression estimates with fixed effects, we confirm that our baseline estimates indeed exhibit a close positive association with TFP calculated based on STAN data.

4.2 Relative Convergence

The summary statistics so far reveal a great deal of variation in how countries' absolute and comparative advantage evolved between the 1960s and today. To shed further light on whether comparative advantage has gotten stronger or weaker over time, we estimate a convergence specification in the spirit of Barro (1991) and Barro and Sala-i-Martin (1992):

$$\Delta \log \left(T_n^j\right)^{1/\theta} = \beta \text{Initial} \log \left(T_n^j\right)^{1/\theta} + \delta_n + \delta_j + \epsilon_{nj} \tag{8}$$

Unlike the classic cross-country convergence regression, our specification pools countries and sectors. On the left-hand side is the log change in the productivity of sector j in country n. The right-hand side regressor of interest is its beginning-of-period value. All of the specifications include country and sector effects, which affects the interpretation of the coefficient. The country effect captures the average change in productivity across all sectors in each country – the absolute advantage. Thus, β picks up the impact of the initial relative productivity on the relative growth of a sector within a country – the evolution of comparative advantage. In particular, a negative value of β implies that relative to the country-specific average, the most backward sectors grew fastest.

Table 5 reports the results. The first column reports the coefficients for the longest differences: the 1960s to the 2000s, while the second column estimates the specification starting in the 1980s. The following 4 columns carry out the estimation decade-by-decade, 1960s to 1970s, 1970s to 1980s, and so on. Since the length of the time period differs across columns, the coefficients are not directly comparable. To help interpret the coefficients, underneath each one we report the speed of convergence, calculated according to the standard Barro and Sala-i-Martin (1992) formula: $\beta = e^{-\lambda \mathcal{T}} - 1$, where β is the regression coefficient on the initial value of productivity, $\mathcal{T}$ is the number of years between the initial and final period, and λ is the convergence speed. This number gives how much of the initial difference between productivities is expected to disappear in a decade. All of the standard errors are clustered by country, to account for unspecified heteroscedasticity at the country level. All of the results are robust to clustering instead at the sector level, and we do not report those standard errors to conserve space.[15]

Column 1 of the top panel reports the estimates for the long-run convergence in the pooled sample of all countries. The coefficient is negative, implying that there is convergence: within a country, the weakest sectors tend to grow faster. It is highly statistically significant: even with clustering the $t-$statistic is over 13. The speed of convergence implied by this coefficient is 24% per decade. As a benchmark, the classic Barro and Sala-i-Martin (1992) rate of convergence is 2% per year, or 22% per decade, strikingly close to what we find in a very different setting.

[15]If the initial T's tend to be measured with error, it has been noted that the convergence regression of the type estimated here will produce bias in favor of finding convergence (Quah 1993). We ran a number of checks to assess the relevance of this effect in our setting. First, we estimated a number of panel specifications with a variety of interacted fixed effects: country×sector, country×decade, and sector×decade included together in estimation. These additional fixed effects will help control for measurement error that varies mainly at country-sector, country-time, or sector-time level, respectively. We also implemented the Arellano-Bond and Blundell-Bond dynamic panel estimators, that difference the data and use lagged values of T to instrument for current changes in T. All of these alternative estimates actually imply a *faster* speed of convergence than the estimates in Table 5. Second, to check how much measurement error is needed to generate our results, we ran a simulation in which we started with artificial data exhibiting zero convergence across sectors within a country, and added measurement error to the right-hand side variable until the OLS coefficient was equal to the coefficient found in our estimates. It turns out that in order for measurememt error to produce coefficient magnitudes found in the data when the truth is zero convergence, it must be the case that 62% of the cross-sectoral variation in the right-hand side variable is due to measurement error.

The second column estimates the long-difference specification from the 1980s to the 2000s. Once again, the coefficient is negative and highly significant, but it implies a considerably slower rate of convergence, 12.4% per decade. The rest of the columns report the results decade-by-decade. Though there is statistically significant convergence in each decade, it is striking that the speed of convergence trends downward, from nearly 30% from the 1960 to the 1970s, to 16.5% in the most recent period.

In order to assess how the results differ across country groups, Panels B and C report the results for the OECD and the non-OECD subsamples separately. (Note that we do not recalculate subsample-specific frontier productivities, so the frontier is the same across subsamples.) Breaking it down produces slightly faster convergence rates than in the full sample. With the exception of the 1980s to the 2000s long difference, the non-OECD countries are catching up somewhat faster, which is not surprising.

Appendix Tables A4 and A5 report the results of estimating the convergence equation (8) country by country, for the periods starting in the 1960s and the 1980s, respectively. These results should be treated with more caution, as the sample size is at most 19. The columns report the coefficient, the standard error, the number of observations, the R^2, as well as the implied speed of convergence for each country. Starting in the 1960s, there is considerable evidence of convergence in these country-specific estimates. In all countries, the convergence coefficient is negative, and significant at the 10% level or below in 39 out of 51 available countries (76%). The evidence starting in the 1980s is weaker: though the large majority of the coefficients are still negative, only 25 out of 61 countries (41%) are showing statistical significance. In addition, most of the countries with a significant coefficient are actually the OECD. Thus, consistent with the pooled results that show a slowdown in convergence starting in the 1980s, these results are less striking than those starting in the 1960s.

All in all, our results provide remarkably robust evidence of relative convergence: in all time periods and broad sets of countries we consider, relatively weak sectors grow faster, with sensible rates of convergence. This implies that Ricardian comparative advantage is getting weaker, at least when measured at the level of broad manufacturing sectors.

4.3 Simple Heuristics: What is Driving These Findings?

What kinds of basic patterns in the data are driving these results? Though our estimation procedure is based on a theoretically-founded gravity equation and a variety of data sources, and thus is fully internally consistent with the underlying conceptual framework, it would be reassuring if we could show some simple heuristic relationships in the data that are consistent with weakening comparative advantage. We can build intuition as follows: in a simpler model with 2 tradeable and 1 nontradeable sectors, Yi and Zhang (2010) show analytically that all else equal, a comparative advantage sector has a smaller share of imports in total domestic absorption $1-\pi_{nn}^{j}$ than a comparative disadvantage sector. As a country's comparative advantage in sector j weakens, the import share rises in that sector. This is intuitive: when a country becomes *relatively* less productive in a sector, it starts importing more.

Thus, weakening comparative advantage should manifest itself in a negative relationship between the initial period import share and the subsequent change in the import share. Sectors within a country with the lowest initial import share $(1-\pi_{nn}^{j})$ should see that import share rise. These are the sectors with the strongest comparative advantage at the beginning of the period. Correspondingly, sectors with the highest initial import share should see their import share drop as they catch up in productivity faster.

Figure 3(a) presents this scatterplot, pooling sectors and countries. The negative relationship is remarkably pronounced: the slope coefficient in the simple bivariate regression is -0.39 with a t-statistic of 14 and an R^2 of 17%. Note that a significant share of the observations – those below zero on the y-axis – have seen their import share actually fall between the 1960s and today. These declines in import shares would be highly puzzling over the period during which trade costs fell and global trade volumes rose dramatically. A strengthening of comparative advantage in those sectors provides a plausible explanation: countries are getting relatively better in those industries, and thus they need to import less.

This negative relationship would not necessarily be evidence of relative convergence in the T's if, for instance, trade costs d_{ni}^{j} fell disproportionally more in sectors in which countries had higher initial import shares. To check for this possibility, Figure 3(b) plots the change in the import-weighted average trade costs in sector j and country n against the initial import share – the same x-axis variable as in the previous figure. There is virtually no relationship between initial import share and subsequent changes in import costs: the slope coefficient, at 0.021, is not statistically significant from zero, and the R^2 is correspondingly 0.00. Thus, it does not appear that systematically larger reductions in d_{ni}^{j} in the initial comparative disadvantage sectors were primarily responsible for the pattern in Figure 3(a).

We can also illustrate the basic patterns in the data using some examples of individual countries. Figure 4 presents the scatterplots of tradable-sector productivities in the 1960s and the 2000s for four countries. The x-axis labels sectors by distance of a sector to the global productivity frontier in the 1960s, so that points further to the left are the comparative advantage sectors in the 1960s – they are closest to the global frontier. The two countries in the top panel – Norway and Korea – showed a pronounced weakening of comparative advantage according to our estimates. They exhibit a clear pattern of faster catch-up to the global frontier in sectors that were the least productive in the 1960s. Italy does not exhibit much of a weakening of comparative advantage: while there is productivity catch-up on average, there is no systematic relationship between initial distance to frontier and subsequent productivity growth in a sector. Finally, South Africa shows a strengthening of comparative advantage: sectors that were initially strongest also grew faster in the subsequent decades.

5 Welfare Analysis

This section computes the welfare impact of changes in comparative advantage documented in the previous section. In order to do this, we solve the full model laid out in Section 2 for a variety of values of technology parameters.

5.1 Benchmark Results and Model Fit

The baseline corresponds to the actual values of T_n^j estimated for the 2000s. Before running the counterfactual experiments, we assess the fit of the baseline model in a number of dimensions. The values of technology parameters are estimated based on the gravity relationship in sectoral trade flows and actual factor endowments, thus the model fits bilateral sector-level trade flows as well as the least-squares gravity relationship can deliver. A more important question is whether the levels of factor prices – w and r – implied by the model are close to the values from the data used in calculating technology parameters. Table 6 compares w's and r's in the model and in the data for 2000s.[16] It is clear that the two are very close: the means and the medians match up quite well, and the correlation between model and data wages is 0.99. The correlation in r's is slightly lower, but still quite high at 0.92.

Another metric by which to evaluate the model is overall trade flows. Though the model is based on matching bilateral sector-level trade flows, it may be that aggregating across different sectors and adding a nontradeable sector leads to biases when it comes to overall trade openness. The bottom panel compares manufacturing imports as a share of GDP in the model to the data.[17] We can see that the averages are extremely close, with both means and medians in the model and the data at roughly 20-22%. The correlation is not perfect, but very high at 0.74. Figure 5 presents the comparision of the three variables between the model and the data graphically.

5.2 Single-Country Counterfactual

We begin by evaluating the impact of each country's changes in comparative advantage on its own welfare. In order to do this, we solve the model while keeping comparative advantage fixed

[16]Comparisons based on earlier decades deliver nearly identical results.

[17]The data on manufacturing imports as a share of GDP come from the World Bank's World Development Indicators.

to the 1960s for one country at a time, and record the change in welfare for that country in the counterfactual relative to the baseline. This counterfactual exercise assumes that between the 1960s and today, each country's sectoral productivities relative to the world frontier grew at their geometric average rate, but comparative advantage remained the same as it was in the 1960s. Precisely, the counterfactual T's are calculated as:

$$\frac{\left(T_n^j\right)_{\text{counterfactual}}}{\left(T_F^j\right)_{2000s}} = \frac{\left(T_n^j\right)_{1960s}}{\left(T_F^j\right)_{1960s}} \times \frac{\left(\prod_{k=1}^{J} (T_n^k/T_F^k)_{2000s}\right)^{\frac{1}{J}}}{\left(\prod_{k=1}^{J} (T_n^k/T_F^k)_{1960s}\right)^{\frac{1}{J}}},$$

where T_F^j is the world frontier in sector j, calculated as in Section 4.

The use of geometric averages has two appealing features. The first is that even though the counterfactual T's are calculated to keep their distance to the frontier, the geometric average of counterfactual T's is equal to the geometric average of the country's actual T's in the 2000s. This ensures that the normalization to the frontier does not induce movements up or down of the average productivity in the country, which would confound the meaning of our counterfactual exercise. The second appealing feature is that this formulation produces identical counterfactual T's whether the experiment is carried out on absolute T's or $T^{1/\theta}$'s, which are the mean productivities. We keep productivity in the nontradeable sector at the benchmark value in all the counterfactual experiments, since our focus is on the welfare impact of changes in comparative advantage.

Table 7 summarizes the results, separating the OECD and the non-OECD countries. The table reports the percentage changes in welfare, for the counterfactual relative to the benchmark. Thus, the positive median values in the first column indicate that on average, welfare would have been higher had comparative advantage not changed since the 1960s. This accords well with what is predicted by theory, given the pronounced weakening of comparative advantage we found in the data in Section 4. However, now we can quantify these effects: for the median OECD country, welfare would have been 1.7% higher had its comparative advantage not weakened. For the non-OECD, the impact very similar, 1.9% at the median.[18]

The second notable aspect of the results is the large dispersion. Among the OECD countries, the standard deviation of welfare changes is 1.8%, while for the non-OECD, it is 2.5 times higher, 5.5%. Correspondingly, the OECD changes range from -0.5% to 5.6%, while for the non-OECD, the range is from -9.3% to 27%. Importantly, among the non-OECD countries, welfare changes range from large negative to large positive, indicating that heterogeneity across countries is first-

[18]A related but distinct question is what is the population-weighted average welfare change, since averaging with population weights in effect assigns equal weights to individuals, rather than countries. It turns out that the population-weighted welfare change in the counterfactual relative to benchmark is about 1%, implying that larger countries tend to have smaller welfare changes.

order.

To cross-check these results and compare magnitudes, the bottom panel of Table 7 reports the same summary statistics for the overall gains from trade compared to autarky for the 2000s in the baseline model. It appears that the welfare impact of the evolution of comparative advantage is on average of the same order of magnitude as the total gains from trade. For the median OECD country, the median gains from trade are 5.2%, while for the non-OECD countries, the median total gains from trade are 4.4%. In addition, there are important differences in the extent of variation of welfare gains from trade compared to welfare changes due to technological changes. In both groups of countries, the gains from trade have a standard deviation of about 3% and a range of about 11%: from a minimum of 1 to a maximum of 12%. For the OECD countries, the range of welfare changes due to technology is much smaller, with a standard deviation of less than 2%, and a range of 6 percentage points. However, for the non-OECD countries, technology changes matter much more: they have a standard deviation of 5.5%, and a range of nearly 40 percentage points. In addition, while gains from trade are – of course – always positive, the welfare impact of technological changes takes on both positive and negative values.

How can we make sense of such a wide variation? Theory predicts that on average, countries experiencing a weakening in comparative advantage should see a reduction in welfare, and countries with a strengthening comparative advantage should be better off. We can verify this by correlating the welfare change implied by the counterfactual exercise to our empirical measures of weakening/strengthening of comparative advantage. Figure 6 presents the results. It plots the change in welfare in the counterfactual relative to the benchmark against the percentage change in the coefficient of variation in a country's $T^{1/\theta}$'s calculated in the previous section. An fall in the coefficient of variation implies that dispersion across sectoral productivities decreased in a country over time – a weakening of comparative advantage. We should expect these countries to on average have higher welfare in the counterfactual that instead fixes comparative advantage to its initial value. Figure 6 confirms this conjecture: there is a pronounced negative relationship between these two variables, with a correlation of -0.5.[19]

5.3 Global Counterfactual

The preceding counterfactual describes the impact of changes in comparative advantage in an individual country on welfare in the country itself. Consistent with the simple intuition gleaned from theory, our empirical finding of weakening comparative advantage also implies that on average, a country would have been better off keeping its 1960s comparative advantage, given the technological change actually observed elsewhere in the world. A complementary, and equally in-

[19]This correlation is virtually unchanged if outlier Indonesia is excluded.

teresting question is what would have happened to all countries had comparative advantage been stuck in the 1960s in every country in the world. Panel A of Table 8 reports the welfare results of this counterfactual. It summarizes the percentage change in welfare that would have resulted had the entire world kept its comparative advantage the same as in the 1960s. Once again, a positive number means that welfare is higher in the counterfactual relative to the benchmark: in this case a country is better off living in the counterfactual world.

On average, while we still find that countries are worse off, these welfare losses are smaller than those in the previous counterfactual, in which only one country's comparative advantage was fixed at the 1960s. The median welfare loss to the OECD is 1.2%, and for the non-OECD 0.6%. The range of outcomes is similar, however. For the non-OECD countries, welfare in the counterfactual ranges from a 9.7% gain to a 22.3% loss. For the OECD, the range of outcomes narrows somewhat.

The preceding two sets of results point to the first-order role of trading partners' evolution of comparative advantage for each country's welfare: the welfare loss from technological change is smaller if everyone's technology is evolving, compared to the case in which only one country is changing its comparative advantage. In the next exercise, we sort out which types of trading partners turn out to be most important for a country's welfare. For instance, it is often suggested that changes in comparative advantage in developing countries can reduce welfare in developed ones (see Samuelson 2004, for a recent example). In order to evaluate this claim, we break up the overall welfare effect into two large groups: that driven by technology changes in the OECD, and in the non-OECD. To do this, we run two additional counterfactual exercises: in the first, we keep the comparative advantage in the OECD countries fixed as in the 1960s, and let the non-OECD countries' comparative advantage evolve as it did in the data. This exercise reveals the welfare changes in all of the countries in the world that are due to the evolution of comparative advantage in the OECD only. In the second counterfactual, we keep the non-OECD comparative advantage fixed to the 1960s instead, and let the OECD technology evolve as it did in the data.

Panels B and C of Table 8 report the results. The patterns are striking: observed changes in OECD comparative advantage tended to hurt the OECD countries, but had virtually no effect on the non-OECD countries. The median impact of OECD technological change on the non-OECD countries is 0.0%, and the range is also tiny, from -0.5% to 0.7%. The same is true of the non-OECD technical change: it tended to lower welfare within that group, and had virtually no impact on the OECD.

Figure 7(a) plots for the OECD countries the welfare changes implied by the evolution of comparative advantage in the OECD only on the y-axis against the total welfare changes from the evolution of comparative advantage in the entire world. Figure 7(b) plots instead the changes in welfare in the OECD due to the non-OECD countries' evolution of comparative advantage.

For ease of interpretation, we add a 45-degree line to both plots. Virtually all of the total welfare change in the OECD is driven by changes in comparative advantage in the OECD itself, as shown in Figure 7(a). By contrast, the non-OECD impact on the OECD is virtually zero for almost all countries. These results imply that while it is true that changes in comparative advantage can lower welfare, for the OECD welfare changes are driven almost exclusively by what happens within that group of countries. These results are due in part to the fact that trade between the OECD countries still accounts for the majority of world trade, and thus the OECD countries are almost always each other's largest trading partners.

Figure 8 repeats the exercise for the non-OECD country group. In Figure 8(a), we plot the non-OECD welfare change in that is due to the OECD comparative advantage changes against the total welfare change. In Figure 8(b), we instead plot the welfare change due to the non-OECD comparative advantage changes. The results are remarkable: among the non-OECD countries, most welfare changes are driven by the non-OECD comparative advantage changes, with virtually no impact from changes in the OECD.

These numbers cannot be explained by the preponderance of trade in this group of countries, since the non-OECD–non-OECD trade is the smallest category of world trade, much lower than the OECD–non-OECD trade. What would be the intuition for these results? Closer inspection of the productivity estimates reveals, not surprisingly, that the frontier productivity in all sectors and decades belongs to the OECD countries. Thus, from the perspective of a non-OECD country trading with the OECD as a group, it essentially always faces trade with the world frontier. While across decades, the particular OECD countries that occupy the frontier position may change, for a typical developing country these changes in comparative advantage in individual OECD countries turn out not to matter much.

5.4 Changes in Comparative Advantage and Trade Volumes

A related aspect of weakening comparative advantage is its impact on trade volumes. Intuition based on simple theory tells us that when comparative advantage weakens, trade volumes should decrease. We confirm this in Table 9, which reports the absolute change in the ratio of imports to GDP in the counterfactual compared to the benchmark. Panel A reports the results for the change in the imports/GDP ratio under the first counterfactual, in which only one country's comparative advantage is kept fixed to the 1960s, while all other countries' sectoral productivities are the same as estimated in the data. For the OECD countries, imports are 1.9 percentage points of GDP higher in the counterfactual compared to the baseline, a proportional increase of about 10% relative to what is observed in the data. For the non-OECD countries, the change is even larger, 4.2 percentage points of GDP, or about a 20% change in trade openness compared to the baseline. Panel B of Table 9 reports the results for the second counterfactual, in which the worldwide relative technology is fixed to the 1960s. Here, the increase is slightly more subdued, 1.8 percentage points of GDP for the OECD, and 2.6 percentage points of GDP for the non-OECD.

6 Conclusion

How does technology evolve over time, and what are the consequences of technological change? In the growth literature, it is widely recognized that economic growth is driven in large part by productivity growth, making it the key force for improvements in welfare. However, when *relative* technology differences are a source of international trade as in the Ricardian world, the welfare impact of technological progress depends on which sectors grow in which countries.

This paper starts by estimating comparative advantage in a sample of some 75 countries, 19 sectors, and 5 decades, 1960s to today. We document a striking pattern in the data: in the world as a whole, comparative advantage is getting weaker over time. This effect is present in all time periods and major country groups: within a country, sectors with the lowest initial relative productivity experience systematically faster productivity growth than sectors with highest initial productivity. This empirical finding opens the door to the theoretical possibility that this type of uneven technological progress can actually reduce welfare in the trading countries. Calibrating the model and solving for the counterfactual scenario in which comparative advantage is instead fixed at its initial-period values, we indeed find that welfare was reduced by weakening comparative advantage. The average impact is large, roughly the same order of magnitude as the total gains from trade for these countries in the 2000s.

In developed countries, the typical worry is that rapid technological catch-up in developing world can lower welfare through this channel. However, we find that nearly all of the welfare impact for the OECD countries comes from changes in comparative advantage within the OECD. Thus, while the negative welfare impact of uneven technological change is very much a feature of the data, for developed countries the culprit is not the poor countries, but rather the rich countries themselves.

The focus of this paper is on measuring how comparative advantage has evolved, and quantifying the welfare impact of this evolution. This exercise leaves open the question of what are the forces driving technological progress and diffusion across countries at the sectoral level. One direction of future research will explore the theoretical mechanisms that could endogenize the patterns that we uncover. The other direction will identify empirically the factors that can account for the evolution of comparative advantage, such as import or export competition, the nature of trading partners, industrial policy, and so on. These two directions are complementary and fruitful avenues for future research.

Appendix A Comparison of Estimated T's with Measured TFP

This Appendix compares the productivity estimates obtained by our procedure and used throughout the paper with estimates of measured TFP that can be obtained directly. Estimating sectoral measured TFP requires data on total output, employment, capital stocks, and intermediate input usage, all in real terms, by sector. This information is only available at sector level and on a consistent basis for many countries through the OECD Structural Analysis (STAN) database. We first compute sectoral capital stocks using data on real investment and the perpetual inventory method, a sectoral equivalent of the procedure described in Section 3.3 for building aggregate capital stocks.[20] We then proceed to compute sector-level measured TFP from data on total output, employment, capital and inputs following equation (6), for all the countries for which required data are available. The set of countries and sectors for which this measured TFP can be computed is not large. There are only 12 countries with all the required data in at least some sectors: Austria, Belgium, Czech Republic, Denmark, Finland, France, Greece, Italy, Norway, Slovenia, Sweden, and United States.[21] The data are in principle available for the period 1970-2008, though in practice earlier years are often not available in individual countries.[22]

It is now well understood that differences in trade openness across sectors will affect measured TFP systematically (see Finicelli et al. 2009a, and Section 3.2). To go from measured TFP to true underlying TFP, we apply the Finicelli et al. (2009a) correction specified for the U.S. in equation (7) to all countries and sectors.

We then correlate the TFP values estimated based on STAN with the T's from our baseline procedure. We present the comparison for the 2000s, as the latest time period has the largest number of observations, and the measures of capital stocks are also more reliable. Panel A of Table A9 reports, for each sector, the Spearman rank correlation between the two measures. These tend to be high: both the mean and median correlations across sectors are 0.6. The last column reports the number of countries for which STAN-based TFP is available in each sector. We can see that most sectors only have information for less than 10 countries. To make more efficient use of the data, we next pool the sectors and examine the correlation between the two productivity measures in a regression framework:

$$\log \text{TFP-STAN}_n^j = \beta \log \left(T_n^j\right)^{1/\theta} + \delta_n + \delta_j + \epsilon_{nj},$$

[20] Though the STAN database contains a variable for sectoral capital stock, it is only available for 6 countries.

[21] In practice, the main bottleneck appears to be data on investment, and therefore capital stocks.

[22] An alternative source of sector-level productivity estimates is the Groningen Growth and Development Centre Productivity Level Database (http://www.ggdc.net/databases/levels.htm). These data are available only at a single point in time, 1997. The database reports levels of multifactor productivity relative to the U.S. for 12 manufacturing sectors and 19 developed countries. We repeated the analysis below using the Groningen data instead. Though the sector-level correlations were somewhat lower than what is reported for STAN, the coefficients from the fixed effects regression were more significant, and the Partial R^2 higher than for STAN.

where TFP-STAN$_n^j$ is the TFP as implied by the STAN data, and T_n^j is as defined in the rest of the paper. The specification includes both country and sector effects, and thus the average productivity levels in individual countries and sectors are netted out. Panel B of Table A9 reports the results. The first column reports the simple bivariate regression of the two measures. The coefficient is close to 1 and highly statistically significant. The correlation between the two variables is 0.341. The second column adds sector effects. The coefficient remains statistically significant at the 1% level, and the partial correlation, obtained after netting the sector effects from both measures of productivity, is much higher at 0.583. Not surprisingly, this is essentially the same as the average correlation within an individual sector in Panel A. Finally, column (3) includes both sector and country effects. The coefficient of interest is significant at the 6.1% level, with the t-statistic of 1.89. With country and sector fixed effects, the overall R^2 is about 0.85. Given that, it is remarkable that the partial correlation between the two measures, after controlling for both country and sector effects is 0.175. Thus, even after netting out all the sector and country effects, the association between these two variables is close and statistically significant.

We can also check how our productivity estimates compare with the aggregate measured TFPs derived from the Penn World Tables using standard Solow residual methods (see, among many others, Hall and Jones 1999, Bernanke and Gürkaynak 2001, Caselli 2005). Appendix Figure A1 plots for the 2000s the Solow residual relative to the US on the x-axis against the output-weighted average of $T^{1/\theta}$ relative to the U.S. on the y-axis, along with the 45-degree line. The aggregate productivity implied by our sector-level estimates of T aligns quite closely with the conventional TFP measure. The correlation between the two in our sample of 75 countries is 0.73. The average levels are quite similar as well. While the median TFP relative to the U.S. according to the Solow residual is 0.50, in our estimates it is equal to 0.55.

We conclude from these exercises that our estimation procedure that relies on bilateral trade to measure productivity delivers results that are in line with the more conventional approaches.

Appendix B Solution Algorithm

Given $\{L_n, K_n, \{T_n^j\}_{j=1}^{J+1}, \xi_n\}_{n=1}^N$, $\{\varepsilon, \alpha_j, \theta, \beta_j, \{\gamma_{k,j}\}_{k=1}^{J+1}, \{d_{ni}^j\}_{N\times N}\}_{j=1}^{J+1}$, and η, we compute the competitive equilibrium of the model as follows.

1. Guess $\{w_n, r_n\}_{n=1}^N$.

    - Compute prices from the following equations:

$$c_n^j = \left(w_n^{\alpha_j} r_n^{1-\alpha_j}\right)^{\beta_j} \left(\prod_{k=1}^{J+1} \left(p_n^k\right)^{\gamma_{k,j}}\right)^{1-\beta_j} \quad \text{for all } n \in \{1, ..., N\} \text{ and } j \in \{1, ..., J+1\},$$

$$\Phi_n^j = \sum_{i=1}^{N} T_i^j \left(c_i^j d_{ni}^j\right)^{-\theta} \text{ for all } n \in \{1, ..., N\} \text{ and } j \in \{1, ..., J\},$$

$$\Phi_n^{J+1} = T_n^{J+1} \left(c_n^{J+1}\right)^{-\theta} \text{ for all } n \in \{1, ..., N\},$$

$$p_n^j = \Gamma \left(\Phi_n^j\right)^{-\frac{1}{\theta}} \text{ for all } n \in \{1, ..., N\} \text{ and } j \in \{1, ..., J+1\},$$

$$P_n = B_n \left(\sum_{j=1}^{J} \omega_j (p_n^j)^{1-\eta}\right)^{\frac{1}{1-\eta}\xi_n} (p_n^{J+1})^{1-\xi_n} \text{ for all } n \in \{1, ..., N\}.$$

- Compute the final demand as follows: for any country n,

$$Y_n^j = \xi_n \frac{w_n L_n + r_n K_n}{p_n^j} \frac{\omega_j (p_n^j)^{1-\eta}}{\sum_{k=1}^{J} \omega_k (p_n^k)^{1-\eta}}, \text{ for all } j = \{1, .., J\},$$

$$Y_n^{J+1} = (1 - \xi_n) \frac{w_n L_n + r_n K_n}{p_n^{J+1}}.$$

- Compute the trade shares π_{ni}^j as follows:

$$\pi_{ni}^j = \frac{T_i^j \left(c_i^j d_{ni}^j\right)^{-\theta}}{\Phi_n^j}.$$

- Compute the total demand as follows: for any country n and any sector j

$$p_n^j Y_n^j + \sum_{k=1}^{J} (\sum_{i=1}^{N} Q_i^k p_i^k \pi_{in}^k)(1 - \beta_k)\gamma_{j,k} + Q_n^{J+1} p_n^{J+1} (1 - \beta_{J+1})\gamma_{j,J+1} = p_n^j Q_n^j.$$

- Compute the factor allocations across sectors as follows: for any country n,

$$\sum_{i=1}^{N} p_i^j Q_i^j \pi_{in}^j = \frac{w_n L_n^j}{\alpha_j \beta_j} = \frac{r_n K_n^j}{(1 - \alpha_j)\beta_j}, \text{ for all } j = \{1, .., J\},$$

$$p_n^{J+1} Q_n^{J+1} = \frac{w_n L_n^{J+1}}{\alpha_{J+1} \beta_{J+1}} = \frac{r_n K_n^{J+1}}{(1 - \alpha_{J+1})\beta_{J+1}}.$$

2. Update $\{w_n', r_n'\}_{n=1}^{N}$ with the feasibility conditions for factors: for any n,

$$\sum_{j=1}^{J+1} L_n^j = L_n, \quad \sum_{j=1}^{J+1} K_n^j = K_n.$$

3. Repeat the above procedures until $\{w_n', r_n'\}_{n=1}^{N}$ is close enough to $\{w_n, r_n\}_{n=1}^{N}$.

References

Barro, Robert J., "Economic Growth in a Cross Section of Countries," *Quarterly Journal of Economics*, May 1991, *106* (2), 407–443.

—— **and Xavier Sala-i-Martin**, "Convergence," *Journal of Political Economy*, April 1992, *100* (2), 223–251.

Bartelsman, Eric J. and Wayne Gray, "The NBER Manufacturing Productivity Database," October 1996. NBER Technical Working Paper 205.

Bernanke, Ben and Refet Gürkaynak, "Is Growth Exogenous? Taking Mankiw, Romer, and Weil Seriously," *NBER Macroeconomics Annual*, 2001, *16*, 11–57.

Bernard, Andrew B. and Charles I. Jones, "Technology and Convergence," *Economic Journal*, July 1996a, *106*, 1037–1044.

—— **and** —— , "Comparing Apples to Oranges: Productivity Convergence and Measurement Across Industries and Countries," *American Economic Review*, December 1996b, *86*, 1216–1238.

Berthelon, Matias and Caroline Freund, "On the Conservation of Distance in International Trade," *Journal of International Economics*, July 2008, *75* (2), 310–320.

Bloom, Nicholas, Mirko Draca, and John Van Reenen, "Trade Induced Technical Change: The Impact of Chinese Imports on Innovation, Diffusion and Productivity," January 2011. NBER WP 16717.

—— **, Paul Romer, and John Van Reenen**, "A Trapped Factors Model of Innovation," October 2010. Mimeo, Stanford University and LSE.

Brezis, Elise S., Paul R. Krugman, and Daniel Tsiddon, "Leapfrogging in International Competition: A Theory of Cycles in National Technological Leadership," *American Economic Review*, December 1993, *83* (5), 1211019.

Burstein, Ariel and Jonathan Vogel, "Globalization, Technology, and the Skill Premium," October 2009. mimeo, UCLA and Columbia University.

Caliendo, Lorenzo and Fernando Parro, "Estimates of the Trade and Welfare Effects of NAFTA," January 2010. mimeo, University of Chicago.

Carrère, Céline, Olivier Cadot, and Vanessa Strauss-Kahn, "Export Diversification: What's behind the Hump?," November 2009. Forthcoming, *Review of Economics and Statistics.*

Caselli, Francesco, "Accounting for Cross-Country Income Differences," in Steven Durlauf Philippe Aghion, ed., *Handbook of Economic Growth*, Vol. 1, Elsevier-North Holland, 2005, chapter 9, pp. 679–741.

—— **and James Feyrer**, "The Marginal Product of Capital," *Quarterly Journal of Economics*, May 2007, *122* (2), 535–568.

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

Costinot, Arnaud, Dave Donaldson, and Ivana Komunjer, "What Goods Do Countries Trade? A Quantitative Exploration of Ricardo's Ideas," April 2011. Forthcoming *Review of Economic Studies.*

di Giovanni, Julian and Andrei A. Levchenko, "Putting the Parts Together: Trade, Vertical Linkages, and Business Cycle Comovement," *American Economic Journal: Macroeconomics*, April 2010, *2* (2), 95–124.

Do, Quy-Toan and Andrei A. Levchenko, "Comparative Advantage, Demand for External Finance, and Financial Development," *Journal of Financial Economics*, December 2007, *86* (3).

Eaton, Jonathan and Samuel Kortum, "Technology, Geography, and Trade," *Econometrica*, September 2002, *70* (5), 1741–1779.

___ , ___ , Brent Neiman, and John Romalis, "Trade and the Global Recession," July 2010. mimeo, Penn State University and University of Chicago.

Finicelli, Andrea, Patrizio Pagano, and Massimo Sbracia, "Ricardian Selection," October 2009a. Bank of Italy *Temi di Discussione* (Working Paper) No. 728.

___ , ___ , and ___ , "Trade-revealed TFP," October 2009b. Bank of Italy *Temi di Discussione* (Working Paper) No. 729.

Hall, Robert and Charles Jones, "Why Do Some Countries Produce So Much More Output per Worker then Others," *Quarterly Journal of Economics*, 1999, *114*, 83–116.

Hausmann, Ricardo and Bailey Klinger, "The Structure of the Product Space and the Evolution of Comparative Advantage," April 2007. CID Working Paper No. 146.

Heston, Alan, Robert Summers, and Bettina Aten, "Penn World Table Version 6.1," October 2002. Center for International Comparisons at the University of Pennsylvania (CICUP).

Hicks, John, "An Inaugural Lecture," *Oxford Economic Papers*, 1953, *5* (2), 117–135.

Hsieh, Chang-Tai and Ralph Ossa, "A global view of productivity growth in China," February 2011. NBER Working Paper No. 16778.

Hummels, David, Jun Ishii, and Kei-Mu Yi, "The Nature and Growth of Vertical Specialization in World Trade," *Journal of International Economics*, June 2001, *54*, 75–96.

Hymans, Saul H. and Frank P. Stafford, "Divergence, Convergence, and the Gains from Trade," *Review of International Economics*, 1995, *3* (1), 118–123.

Imbs, Jean and Romain Wacziarg, "Stages of Diversification," *American Economic Review*, March 2003, *93* (1), 63–86.

Jones, Ronald, "Technical Progress and Real Income in a Ricardian Trade Model," in Ronald Jones, ed., *Internarional Trade: Essays in Theory*, Amsterdam: North-Holland, 1979.

Keller, Wolfgang, "International Technology Diffusion," *Journal of Economic Literature*, September 2004, *42* (3), 752–782.

Klenow, Peter J. and Andrés Rodríguez-Clare, "The Neoclassical Revival in Growth Economics: Has It Gone Too Far?," *NBER Macroeconomics Annual*, 1997, *12*, 73–103.

Krugman, Paul, "A Model of Innovation, Technology Transfer, and the World Distribution of Income," *Journal of Political Economy*, April 1979, *87* (2), 253–66.

___ , "The narrow moving band, the Dutch disease, and the competitive consequences of Mrs. Thatcher: Notes on trade in the presence of dynamic scale economies," *Journal of Development Economics*, 1987, *27* (1-2), 41–55.

Proudman, James and Stephen Redding, "Evolving Patterns of International Trade," *Review of International Economics*, 2000, *8* (3), 373–396.

Quah, Danny, "Galton's Fallacy and Tests of the Convergence Hypothesis," *Scandinavian Journal of Economics*, December 1993, *95* (4), 427–443.

Rose, Andrew K., "Do We Really Know That the WTO Increases Trade?," *American Economic Review*, March 2004, *94* (1), 98–114.

Samuelson, Paul A., "Where Ricardo and Mill Rebut and Confirm Arguments of Mainstream Economists Supporting Globalization," *Journal of Economic Perspectives*, 2004, *18* (3), 135–146.

Shikher, Serge, "Putting industries into the Eaton-Kortum model," July 2004. Forthcoming, *Journal of International Trade and Economic Development*.

——, "Accounting for International Trade," August 2005. mimeo, Suffolk University.

——, "Capital, technology, and specialization in the neoclassical model," *Journal of International Economics*, March 2011, *83* (2), 229–242.

Simonovska, Ina and Michael E. Waugh, "The Elasticity of Trade: Estimates and Evidence," December 2010. Mimeo, UC Davis and NYU.

Waugh, Michael, "International Trade and Income Differences," *American Economic Review*, December 2010, *100* (5), 2093–2124.

Yi, Kei-Mu and Jing Zhang, "Structural Change in an Open Economy," April 2010. Mimeo, Federal Reserve Bank of Philadelphia and University of Michigan.

Young, Alwyn, "Learning by Doing and the Dynamic Effects of International Trade," *Quarterly Journal of Economics*, May 1991, *106* (2), 369–405.

Table 1. Summary Statistics

	OECD			Non-OECD		
	Mean $T^{1/\theta}$	Top2/bottom2 $T^{1/\theta}$	Countries	Mean $T^{1/\theta}$	Top2/bottom2 $T^{1/\theta}$	Countries
1960s	0.651	1.502	21	0.453	2.066	33
1970s	0.692	1.434	21	0.471	1.775	37
1980s	0.776	1.412	22	0.509	1.922	42
1990s	0.808	1.395	22	0.378	2.136	53
2000s	0.838	1.394	22	0.410	2.088	53

Notes: This table reports the summary statistics for the average productivity relative to the frontier (mean $T^{1/\theta}$), the relative productivity of the two most productive tradeable sectors relative to the 2 least productive ones (top2/bottom2 $T^{1/\theta}$), as well as the number of countries for which data are available. The samples are split by decade and into OECD and non-OECD groups.

Table 2. Average Convergence: Fastest and Slowest Countries

Since 1960s		Since 1980s	
Top 10: Fastest Converging Countries		Top 10: Fastest Converging Countries	
Iceland	0.618	Portugal	0.373
Norway	0.615	Greece	0.364
Korea, Rep.	0.566	Ireland	0.315
Ireland	0.525	Norway	0.258
Netherlands	0.449	Iceland	0.240
Finland	0.445	Korea, Rep.	0.240
Israel	0.384	Belgium-Luxembourg	0.182
Greece	0.382	Mauritius	0.162
Portugal	0.347	United Kingdom	0.159
Germany	0.337	Finland	0.138
Bottom 10: Slowest Converging Countries		Bottom 10: Slowest Converging Countries	
Malaysia	-0.163	Senegal	-0.226
Philippines	-0.166	Argentina	-0.236
Canada	-0.183	Brazil	-0.237
Turkey	-0.259	Peru	-0.270
Thailand	-0.271	India	-0.332
Venezuela, RB	-0.276	Iran, Islamic Rep.	-0.348
Honduras	-0.337	Venezuela, RB	-0.366
India	-0.358	Ethiopia	-0.395
Egypt, Arab Rep.	-0.372	Egypt, Arab Rep.	-0.405
Sri Lanka	-0.419	Honduras	-0.428

Notes: This table reports the 10 fastest and 10 slowest converging countries since the 1960s (left panel) and the 1980s (right panel), measured by the percent change in the mean absolute distance to the frontier across all tradeable sectors.

Table 3. Relative Convergence: Fastest and Slowest Countries

Since 1960s		Since 1980s	
Top 10: Fastest Converging Countries		Top 10: Fastest Converging Countries	
Norway	-0.654	Norway	-0.534
Indonesia	-0.396	Sweden	-0.379
Finland	-0.379	Greece	-0.264
Sweden	-0.343	Denmark	-0.231
Spain	-0.333	Iceland	-0.199
Korea, Rep.	-0.327	Finland	-0.185
Denmark	-0.299	Spain	-0.172
Belgium-Luxembourg	-0.290	Chile	-0.142
Iceland	-0.286	Germany	-0.133
Ireland	-0.271	Costa Rica	-0.111
Bottom 10: Slowest Converging Countries		Bottom 10: Slowest Converging Countries	
India	0.132	Trinidad and Tobago	0.301
Kenya	0.154	Saudi Arabia	0.308
Honduras	0.185	Italy	0.317
Thailand	0.260	El Salvador	0.352
Egypt, Arab Rep.	0.300	Canada	0.352
South Africa	0.315	Australia	0.419
Ghana	0.353	Venezuela, RB	0.584
Japan	0.448	Egypt, Arab Rep.	0.761
Canada	0.485	Iran, Islamic Rep.	0.785
Sri Lanka	0.744	Japan	0.880

Notes: This table reports the 10 fastest and 10 slowest converging countries since the 1960s (left panel) and the 1980s (right panel), measured by the percent change in the coefficient of variation across tradeable sectors in the distance to the frontier.

Table 4. Correlations Between Convergence Measures, Per Capita Income Growth, and Changes in Openness

Since the 1960s	Pct Chg in Average Abs. Distance	Pct Chg in Coeff. Var. of $T^{1/\theta}$	Pct Chg in Real Per Capita Income	Pct Chg in Trade Openness
Pct Chg in Average Abs. Distance	*0.270*			
Pct Chg in Coeff. Var. of $T^{1/\theta}$	-0.642	*0.263*		
Pct Chg in Real Per Capita Income	0.238	-0.140	*1.483*	
Pct Chg in Trade Openness	-0.293	-0.074	0.303	*0.981*

Since the 1980s	Average Abs. Distance	Pct Chg in Coeff. Var. of $T^{1/\theta}$	Pct Chg in Real Per Capita Income	Pct Chg in Trade Openness
Pct Chg in Average Abs. Distance	*0.195*			
Pct Chg in Coeff. Var. of $T^{1/\theta}$	-0.608	*0.263*		
Pct Chg in Real Per Capita Income	0.260	-0.052	*0.545*	
Pct Chg in Trade Openness	-0.331	0.048	0.121	*0.504*

Notes: This table reports the correlation coefficients (off-diagonal elements), and standard deviations (diagonal elements, in italics) between the measure of average convergence (Pct Chg in Average Abs. Distance), relative convergence (Pct Chg in Coeff. Var. of T), real PPP-adjusted per capita income, and overall trade openness. The latter two measures come from the Penn World Tables 6.3.

Table 5. Pooled Regression Results

	(1) 1960s to 2000s	(2) 1980s to 2000s	(3) 1960s to 1970s	(4) 1970s to 1980s	(5) 1980s to 1990s	(6) 1990s to 2000s
Dep. Var: Log Change in $T^{1/\theta}$						
	Panel A: All Countries					
Log(Initial$T^{1/\theta}$)	-0.618***	-0.220***	-0.254***	-0.168***	-0.195***	-0.152***
	(0.046)	(0.030)	(0.029)	(0.027)	(0.029)	(0.040)
NB:						
Speed of convergence, per decade	*0.241*	*0.124*	*0.293*	*0.184*	*0.217*	*0.165*
Observations	929	1,122	991	1,074	1,183	1,335
R^2	0.844	0.833	0.851	0.841	0.897	0.863
	Panel B: OECD					
Log(Initial$T^{1/\theta}$)	-0.723***	-0.414***	-0.269***	-0.145***	-0.258***	-0.174**
	(0.092)	(0.063)	(0.042)	(0.036)	(0.048)	(0.074)
NB:						
Speed of convergence, per decade	*0.321*	*0.267*	*0.313*	*0.157*	*0.298*	*0.191*
Observations	393	405	396	394	407	410
R^2	0.860	0.847	0.874	0.839	0.799	0.834
	Panel C: non-OECD					
Log(Initial$T^{1/\theta}$)	-0.731***	-0.269***	-0.378***	-0.227***	-0.264***	-0.206***
	(0.056)	(0.046)	(0.041)	(0.040)	(0.042)	(0.054)
NB:						
Speed of convergence, per decade	*0.328*	*0.157*	*0.475*	*0.257*	*0.307*	*0.231*
Observations	536	717	595	680	776	925
R^2	0.851	0.813	0.868	0.853	0.901	0.873
Country FE	yes	yes	yes	yes	yes	yes
Sector FE	yes	yes	yes	yes	yes	yes

Notes: Standard errors clustered at the country level in parentheses; ***: significant at 1%; **: significant at 5%. This table reports the results of regressing the growth of estimated technology parameter $\left(T_n^j\right)^{1/\theta}$ on its initial value over different time periods and subsamples. The speed of convergence, per decade, is reported (in italics) underneath each coefficient estimate.

Table 6. The Fit of the Baseline Model with the Data

	model	data
Wages:		
mean	0.381	0.333
median	0.125	0.145
corr(model, data)	0.987	
Return to capital:		
mean	0.830	0.919
median	0.632	0.698
corr(model, data)	0.918	
Imports/GDP:		
mean	0.222	0.237
median	0.212	0.200
corr(model, data)	0.739	

Notes: This table reports the means and medians of wages relative to the U.S. (top panel); return to capital relative to the U.S. (middle panel), and imports as a share of GDP (bottom panel), in the model and in the data. In the data, Imports/GDP are the manufacturing imports as a share of GDP in the 2000s, sourced from the World Bank's World Development Indicators. Wages and return to capital in the data are calculated as described in Section 3.3.

Table 7. Welfare Gains in the Single-Country Counterfactual Relative to Baseline

	Median	St. Dev.	Min	Max	Countries
OECD	0.017	0.018	-0.005	0.056	22
Non-OECD	0.019	0.055	-0.093	0.270	53
NB: Overall gains from trade					
OECD	0.052	0.032	0.011	0.120	
Non-OECD	0.044	0.029	0.005	0.122	

Notes: This table reports the proportional change in welfare under the counterfactual scenario with respect to the baseline. The counterfactual assumes that for each individual country, comparative advantage remained as it was in the 1960s, while its T's grew at the same country-specific average rate between the 1960s and the 2000s. All other countries' comparative advantage is taken from the data. In the baseline comparative advantage is as it is in the data for the 2000s. The lower panel reports the total gains from trade relative to autarky in the baseline for the 2000s.

Table 8. Welfare Gains in the Global Counterfactual Relative to Baseline

	Median	St. Dev.	Min	Max	Countries
Panel A: CA fixed to 1960s in all countries					
OECD	0.012	0.013	-0.008	0.038	22
Non-OECD	0.006	0.050	-0.097	0.223	53
Panel B: CA fixed to 1960s in OECD countries only					
OECD	0.013	0.014	-0.008	0.041	
Non-OECD	0.000	0.002	-0.005	0.007	
Panel C: CA fixed to 1960s in non-OECD countries only					
OECD	0.000	0.002	-0.002	0.006	
Non-OECD	0.013	0.054	-0.097	0.257	

Notes: This table reports the proportional change in welfare under the counterfactual scenario with respect to the baseline. The counterfactual assumes that in all countries in the world (Panel A), in OECD (Panel B) and the non-OECD (Panel C), comparative advantage remained as it was in the 1960s, while its T's grew at the same country-specific average rate between the 1960s and the 2000s. In the baseline comparative advantage is as it is in the data for the 2000s.

Table 9. Absolute Change in Imports/GDP in the Counterfactuals Relative to Baseline

	Median	St. Dev.	Min	Max	Countries
Panel A: Country-by-country counterfactual					
OECD	0.019	0.039	-0.008	0.128	22
Non-OECD	0.042	0.079	-0.070	0.430	53
Panel B: Global counterfactual					
OECD	0.018	0.015	-0.004	0.048	
Non-OECD	0.026	0.039	-0.042	0.169	

Notes: This table reports the absolute change in imports/GDP under the counterfactual scenarios with respect to the baseline. In Panel A, the counterfactual scenario assumes that a single country's comparative advantage is the same as in the 1960s, and evaluates the impact of this change for that country's trade volumes. In Panel B, the counterfactual scenario assumes that comparative advantage is fixed to the 1960s in every country in the world, and reports the summary statistics for the change in trade volumes in this sample of countries.

Figure 1. Absolute and Relative Convergence, 1960s – 2000s

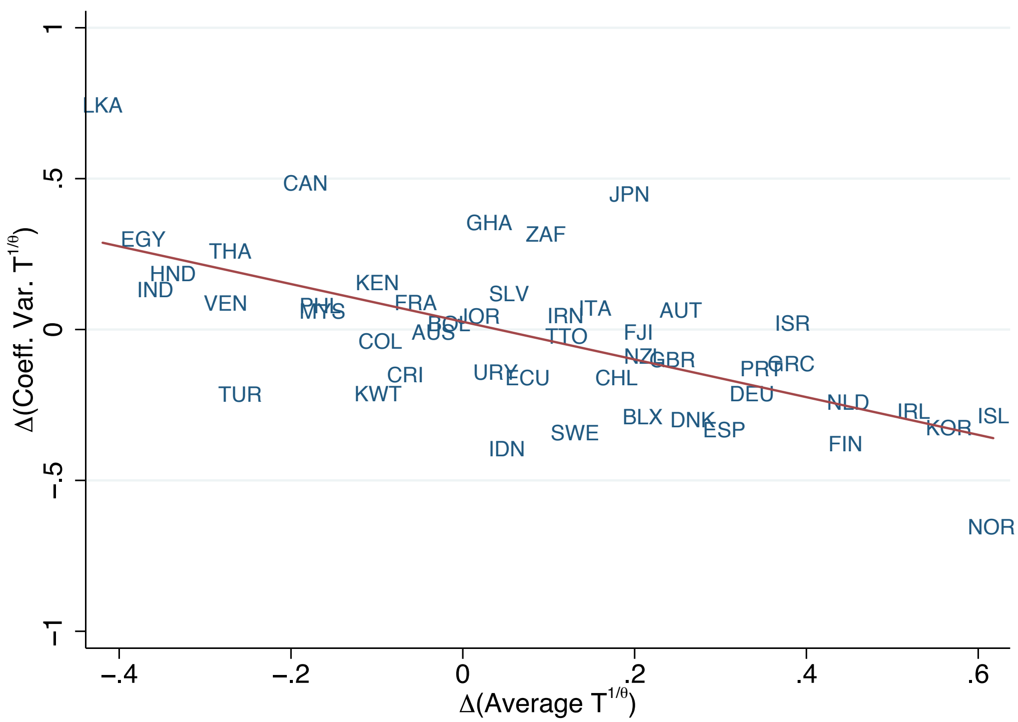

Notes: This figure displays the proportional change in a country's average distance to the world frontier (horizontal axis) against the proportional change in the coefficient of variation in distances to frontier across sectors (vertical axis), along with the least squares fit through the data.

Figure 2. Convergence, Income Growth, and Changes in Trade Openness, 1960s to 2000s

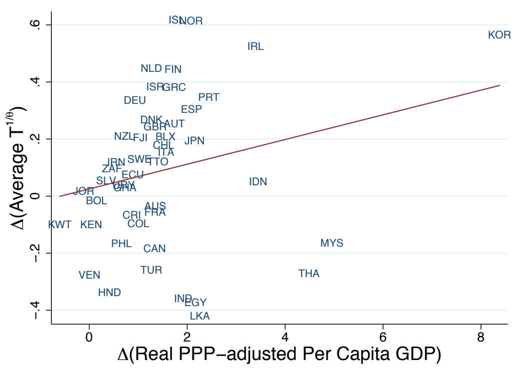

(a) Absolute Convergence and Income Growth

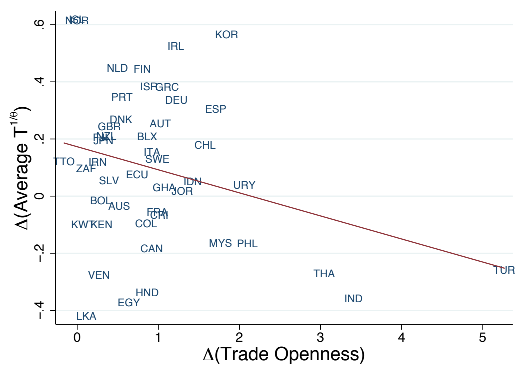

(b) Absolute Convergence and Trade Growth

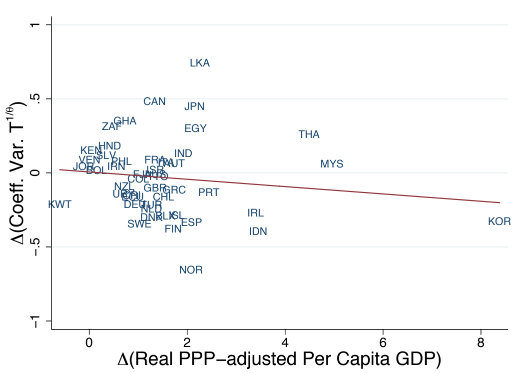

(c) Relative Convergence and Income Growth

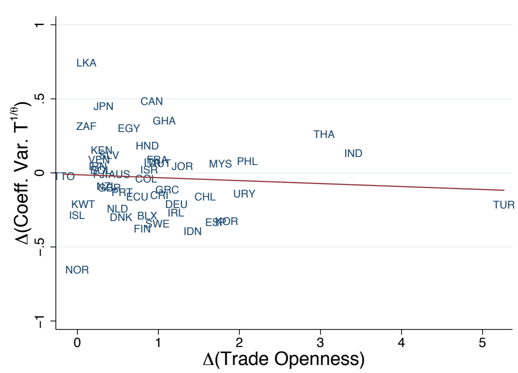

(d) Relative Convergence and Trade Growth

Notes: This figure presents the bivariate plots of absolute (top row) and relative convergence (bottom row), against contemporaneous changes in PPP-adjusted real per capita GDP and changes in trade openness (Imports + Exports)/GDP.

Figure 3. Heuristic Evidence: Initial Import Shares, Changes in Import Shares, and Changes in Trade Costs

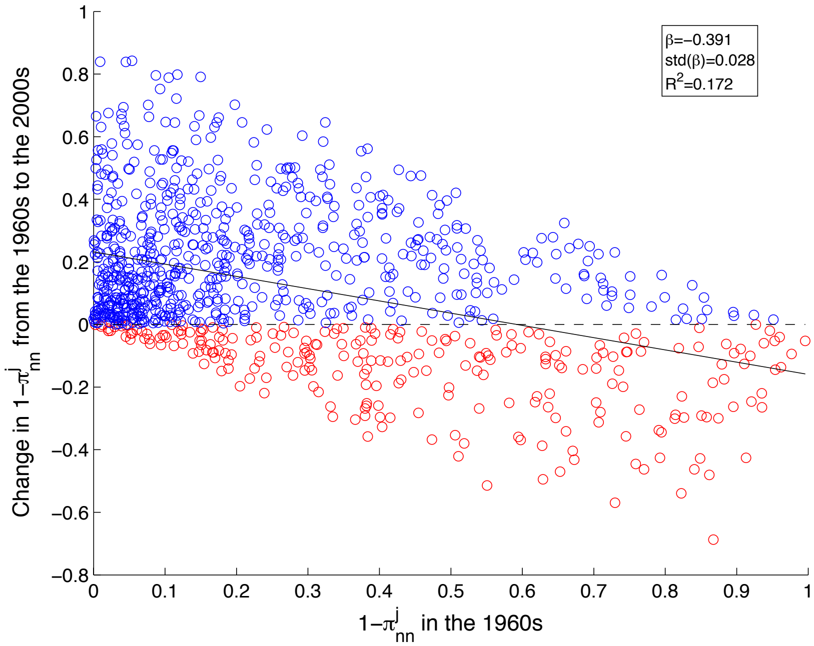

(a) Initial Import Shares and Changes in Import Shares

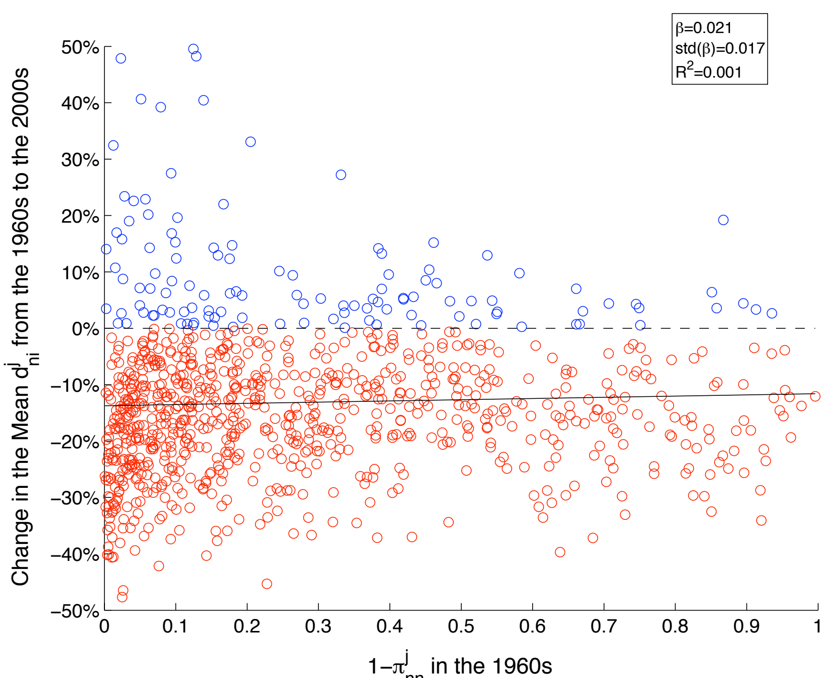

(b) Initial Import Shares and Changes in Trade Costs

Notes: This figure plots the change in the import share between the 1960s and 2000s $\Delta(1-\pi^j_{nn})$ (top panel), and the percentage change in import-weighted average import costs d^j_{ni} between the 1960 and the 2000s (bottom panel), against the import share of sector j in country n in the 1960s on the x-axis. The figure pools country-sectors.

Figure 4. Productivities in Selected Countries

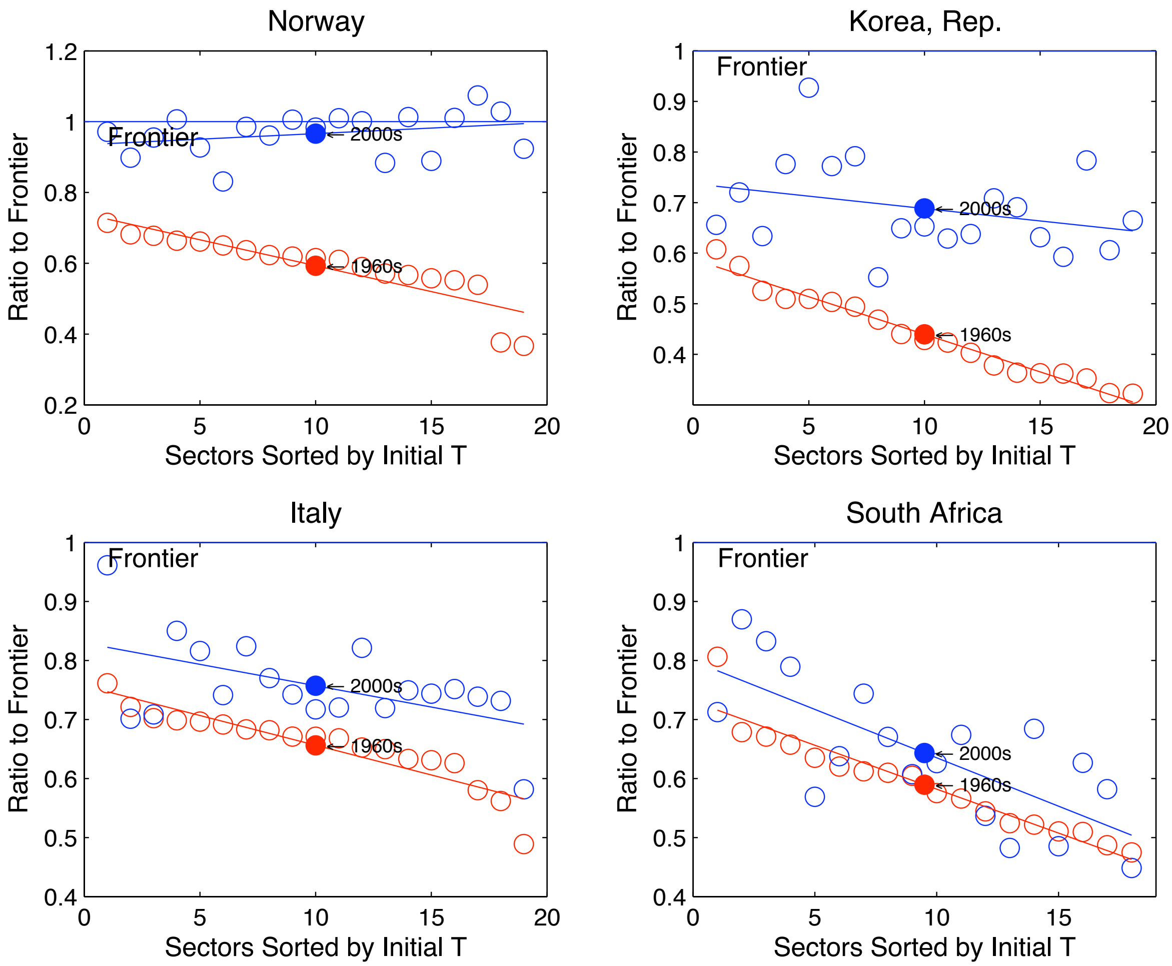

Notes: This figure displays the tradable-sector productivities in selected countries, expressed as a ratio to the global frontier productivity, for the 1960s and the 2000s. The x-axis labels sectors in descending order of the ratio to frontier in the 1960s for each country (so that the most productive sectors in that country are on the left). The lines are OLS fits through the points for each decade. The solid circles are the average of sectoral productivities relative to the global frontier for each decade.

Figure 5. Benchmark Model vs. Data

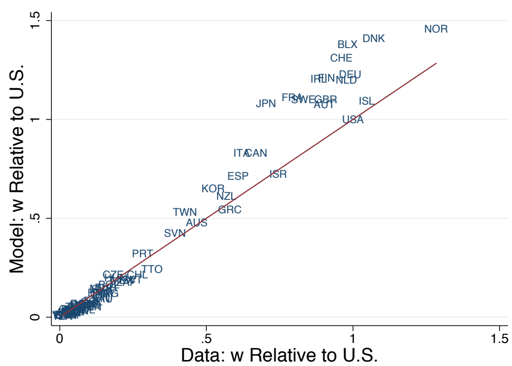

(a) Wages

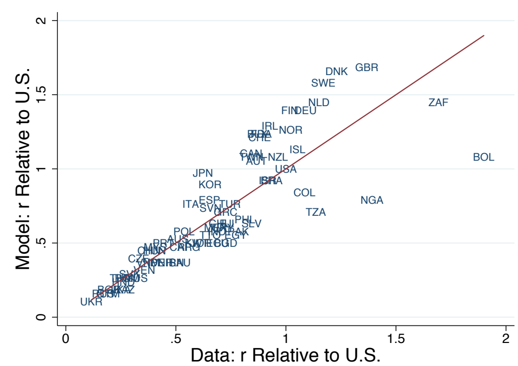

(b) Return to Capital

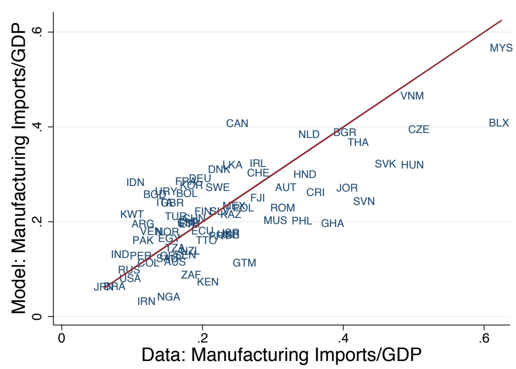

(c) Manufacturing Imports/GDP

Notes: This figure presents the scatterplots of wages, return to capital, and manufacturing imports/GDP, for the model (y-axis) against the data (x-axis). The straight line in each plot is the 45-degree line.

Figure 6. Welfare Changes and Relative Convergence

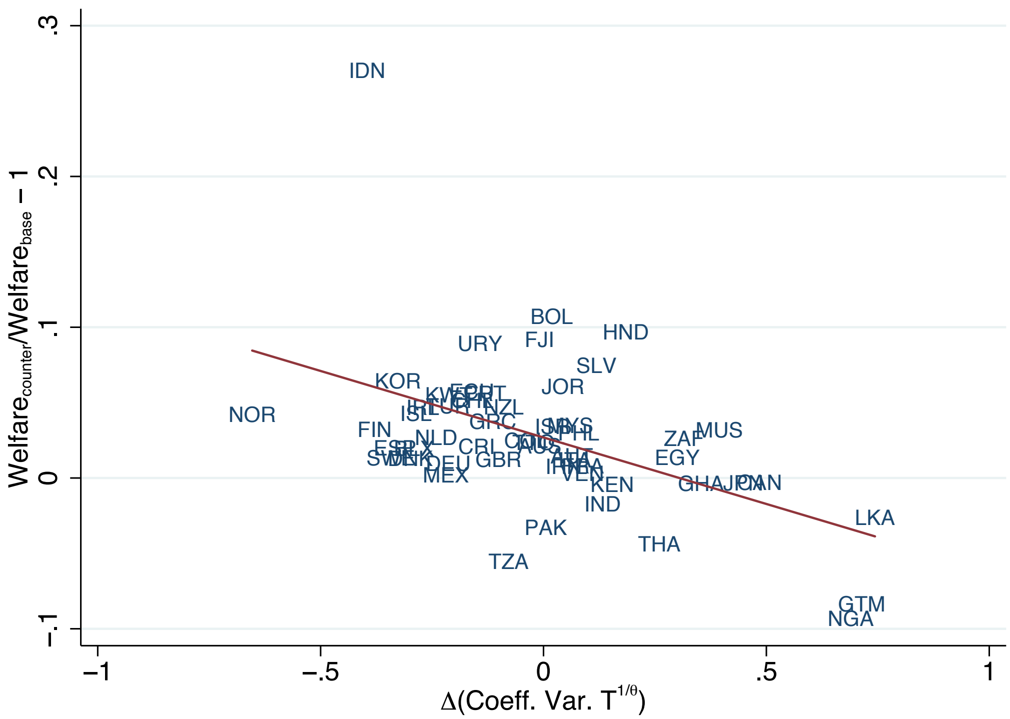

Notes: This figure displays the proportional change in a country's welfare in the counterfactual scenario in which its comparative advantage was fixed at its 1960s value relative to the baseline (y-axis), against the change in the coefficient of variation in the country's $T^{1/\theta}$ between the 1960s and the 2000s (x-axis). A larger value of the x-axis variable implies that comparative advantage has gotten stronger. A negative value implies that comparative advantage has gotten weaker. The solid line is the least-squares fit.

Figure 7. Welfare Changes for OECD Countries

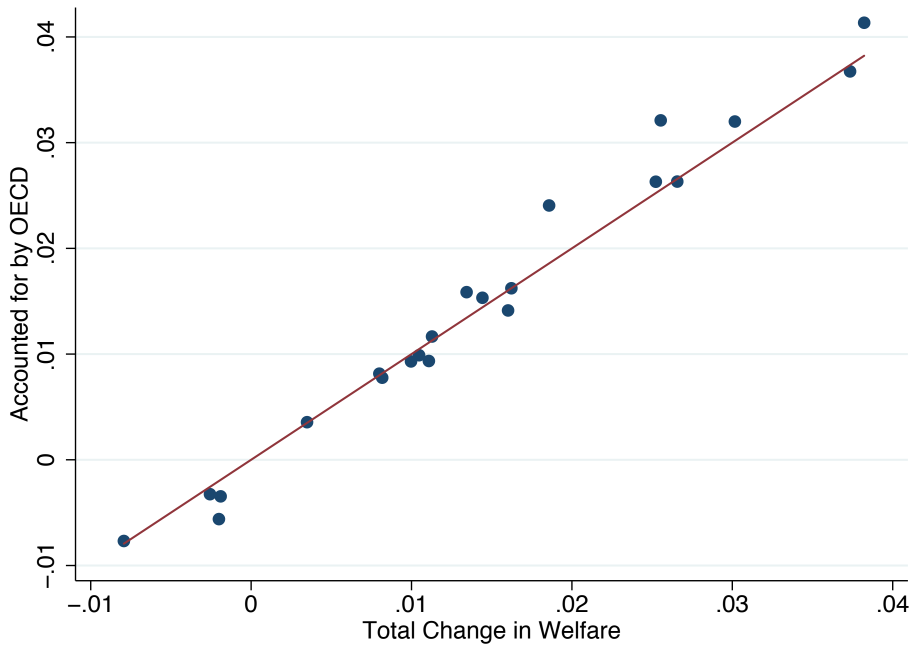

(a) Accounted for by OECD

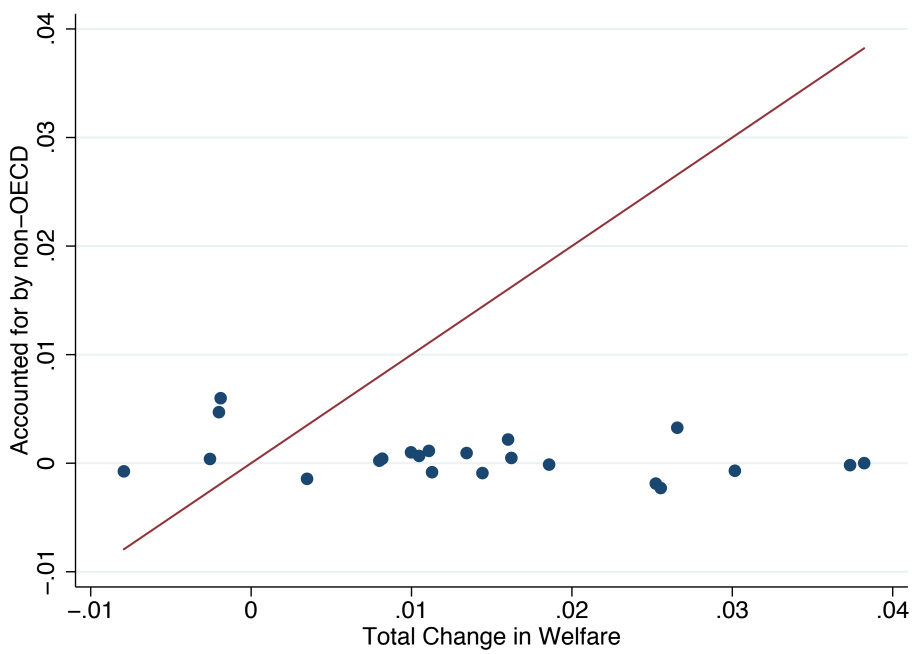

(b) Accounted for by non-OECD

Notes: This figure plots, for the OECD countries, the total welfare change in the global counterfactual on the x-axis against the welfare change due to comparative advantage changes in the OECD only (top panel), and the non-OECD only (bottom panel). The straight line is the 45-degree line.

Figure 8. Welfare Changes for Non-OECD Countries

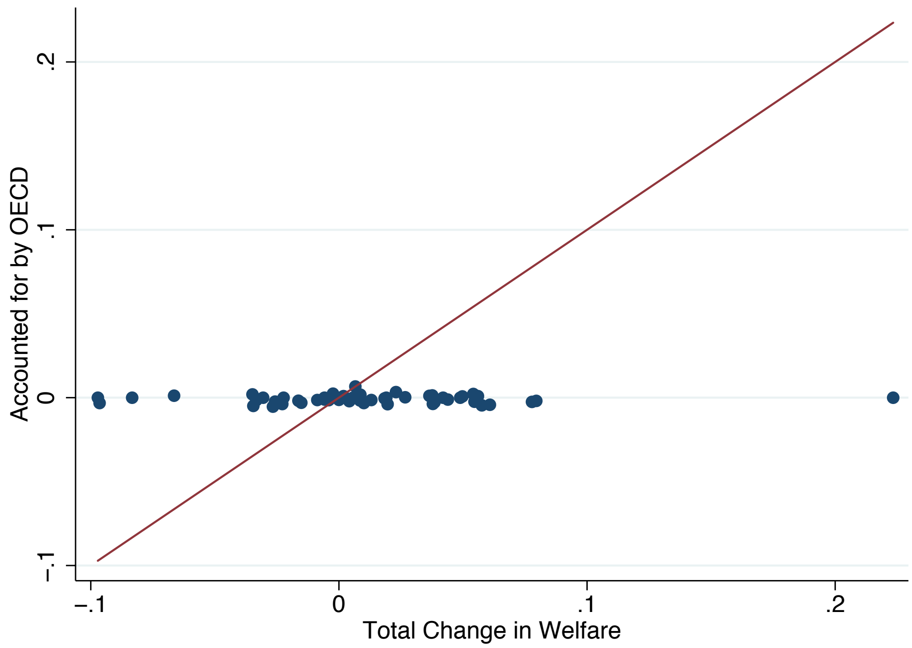

(a) Accounted for by OECD

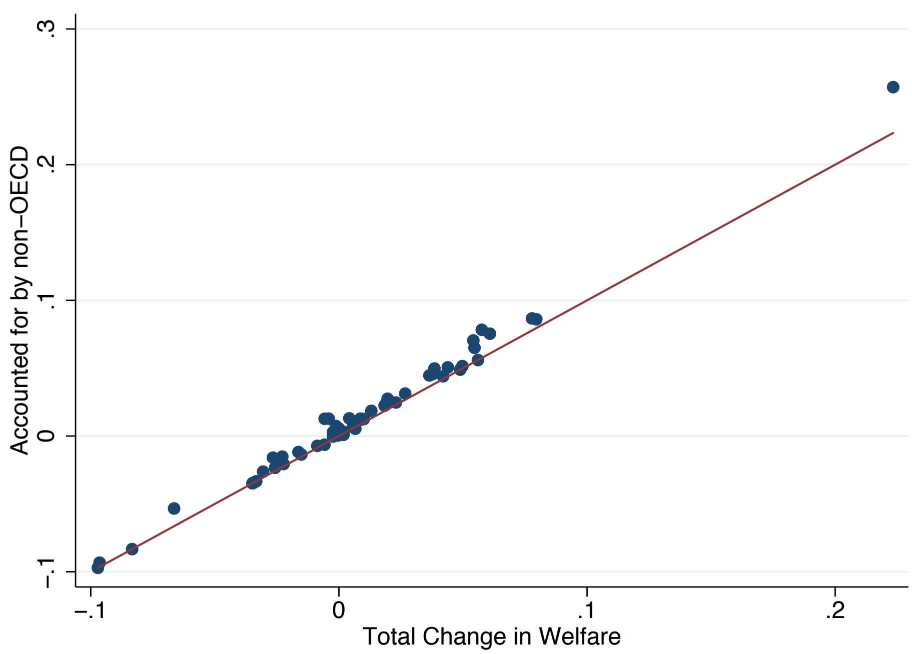

(b) Accounted for by non-OECD

Notes: This figure plots, for the non-OECD countries, the total welfare change in the global counterfactual on the x-axis against the welfare change due to comparative advantage changes in the OECD only (top panel), and the non-OECD only (bottom panel). The straight line is the 45-degree line.

Table A1. Country Coverage

Country	Period	Country	Period
OECD		Non-OECD	
Australia	1960s−2000s	Argentina	1980s−2000s
Austria	1960s−2000s	Bangladesh	1970s−2000s
Belgium-Luxembourg	1960s−2000s	Bolivia	1960s−2000s
Canada	1960s−2000s	Brazil	1980s−2000s
Denmark	1960s−2000s	Bulgaria	1990s−2000s
Finland	1960s−2000s	Chile	1960s−2000s
France	1960s−2000s	China	1970s−2000s
Germany	1960s−2000s	Colombia	1960s−2000s
Greece	1960s−2000s	Costa Rica	1960s−2000s
Iceland	1960s−2000s	Czech Republic	1990s−2000s
Ireland	1960s−2000s	Ecuador	1960s−2000s
Italy	1960s−2000s	Egypt, Arab Rep.	1960s−2000s
Japan	1960s−2000s	El Salvador	1960s−2000s
Netherlands	1960s−2000s	Ethiopia	1980s−2000s
New Zealand	1960s−2000s	Fiji	1960s−2000s
Norway	1960s−2000s	Ghana	1960s−2000s
Portugal	1960s−2000s	Guatemala	1960s−2000s
Spain	1960s−2000s	Honduras	1960s−2000s
Sweden	1960s−2000s	Hungary	1990s−2000s
Switzerland	1980s−2000s	India	1960s−2000s
United Kingdom	1960s−2000s	Indonesia	1960s−2000s
United States	1960s−2000s	Iran, Islamic Rep.	1960s−2000s
		Israel	1960s−2000s
		Jordan	1960s−2000s
		Kazakhstan	1990s−2000s
		Kenya	1960s−2000s
		Korea, Rep.	1960s−2000s
		Kuwait	1960s−2000s
		Malaysia	1960s−2000s
		Mauritius	1960s−2000s
		Mexico	1960s−2000s
		Nigeria	1960s−2000s
		Pakistan	1960s−2000s
		Peru	1980s−2000s
		Philippines	1960s−2000s
		Poland	1990s−2000s
		Romania	1990s−2000s
		Russian Federation	1990s−2000s
		Saudi Arabia	1980s−2000s
		Senegal	1970s−2000s
		Slovak Republic	1990s−2000s
		Slovenia	1990s−2000s
		South Africa	1960s−2000s
		Sri Lanka	1960s−2000s
		Taiwan Province of China	1970s−2000s
		Tanzania	1960s−2000s
		Thailand	1960s−2000s
		Trinidad and Tobago	1960s−2000s
		Turkey	1960s−2000s
		Ukraine	1990s−2000s
		Uruguay	1960s−2000s
		Venezuela, RB	1960s−2000s
		Vietnam	1990s−2000s

Notes: This table reports the countries in the sample and the decades for which data are available for each country.

Table A2. Sectors

ISIC code	Sector Name	α_j	β_j	$\gamma_{J+1,j}$	ω_j
15	Food and Beverages	0.315	0.281	0.303	0.209
16	Tobacco Products	0.264	0.520	0.527	0.010
17	Textiles	0.467	0.371	0.295	0.025
18	Wearing Apparel, Fur	0.493	0.377	0.320	0.089
19	Leather, Leather Products, Footwear	0.485	0.359	0.330	0.014
20	Wood Products (Excl. Furniture)	0.452	0.372	0.288	0.009
21	Paper and Paper Products	0.366	0.344	0.407	0.012
22	Printing and Publishing	0.484	0.469	0.407	0.004
23	Coke, Refined Petroleum Products, Nuclear Fuel	0.244	0.243	0.246	0.092
24	Chemical and Chemical Products	0.308	0.373	0.479	0.008
25	Rubber and Plastics Products	0.385	0.387	0.350	0.014
26	Non-Metallic Mineral Products	0.365	0.459	0.499	0.071
27	Basic Metals	0.381	0.299	0.451	0.002
28	Fabricated Metal Products	0.448	0.398	0.364	0.012
29C	Office, Accounting, Computing, and Other Machinery	0.473	0.390	0.388	0.094
31A	Electrical Machinery, Communication Equipment	0.405	0.380	0.416	0.057
33	Medical, Precision, and Optical Instruments	0.456	0.428	0.441	0.036
34A	Transport Equipment	0.464	0.343	0.286	0.175
36	Furniture and Other Manufacturing	0.460	0.407	0.397	0.065
4A	Nontradeables	0.561	0.651	0.788	
	Mean	0.414	0.393	0.399	0.053
	Min	0.244	0.243	0.246	0.002
	Max	0.561	0.651	0.788	0.209

Notes: This table reports the sectors used in the analysis. The classification corresponds to the ISIC Revision 3 2-digit, aggregated further due to data availability. α_j is the value-added based labor intensity; β_j is the share of value added in total output; $\gamma_{J+1,j}$ is the share of nontradeable inputs in total intermediate inputs; ω_j is the taste parameter for tradeable sector j, estimated using the procedure described in Section 3.2. Variable definitions and sources are described in detail in the text.

Table A3. Zero Trade Observations: Model vs. Data

	ISIC	1960s	1970s	1980s	1990s	2000s
Overall		0.001	0.001	0.002	0.001	0.001
Food and Beverages	15	0.000	0.000	0.000	0.000	0.000
Tobacco Products	16	0.002	0.001	0.001	0.000	0.002
Textiles	17	0.001	0.001	0.005	0.003	0.001
Wearing Apparel Fur	18	0.000	0.005	0.003	0.001	0.000
Leather Leather Products Footwear	19	0.001	0.003	0.008	0.004	0.007
Wood Products (Excl. Furniture)	20	0.002	0.002	0.002	0.002	0.002
Paper and Paper Products	21	0.001	0.001	0.001	0.001	0.000
Printing and Publishing	22	0.000	0.000	0.000	0.000	0.000
Coke Refined Petroleum Products Nuclear Fuel	23	0.000	0.001	0.001	0.000	0.000
Chemical and Chemical Products	24	0.001	0.002	0.002	0.001	0.001
Rubber and Plastics Products	25	0.000	0.001	0.001	0.001	0.000
Non-Metallic Minetal Products	26	0.000	0.000	0.001	0.000	0.001
Basic Metals	27	0.001	0.001	0.002	0.001	0.002
Fabricated Metal Products	28	0.000	0.000	0.001	0.001	0.000
Office Accounting Computing and Other Machinery	29C	0.001	0.001	0.001	0.001	0.000
Electrical Machinery Communication Equipment	31A	0.001	0.001	0.001	0.000	0.000
Medical Precision and Optical Instruments	33	0.001	0.001	0.002	0.001	0.001
Transport Equipment	34A	0.000	0.000	0.001	0.001	0.000
Furniture and Other Manufacturing	36	0.001	0.001	0.001	0.001	0.000

Notes: This table reports the share of global absorption taken up by importer-exporter-sector observations for which actual imports are zero in the data.

Table A4. Country-by-Country Estimates of Relative Convergence, 1960s to 2000s

Country	β	s.e.	Obs.	R^2	Speed of Convergence, by decade
United Kingdom	-0.831***	0.188	19	0.469	0.444
Austria	-0.964**	0.336	19	0.450	0.828
Belgium-Luxembourg	-0.872***	0.188	19	0.660	0.515
Denmark	-1.025***	0.166	19	0.692	–
France	-0.738***	0.198	19	0.343	0.335
Germany	-0.753***	0.138	19	0.527	0.350
Italy	-0.320	0.208	19	0.160	0.096
Netherlands	-0.772***	0.182	19	0.563	0.370
Norway	-1.028***	0.062	19	0.917	–
Sweden	-0.890***	0.178	18	0.544	0.552
Canada	-0.293	0.275	19	0.046	0.087
Japan	-0.831**	0.304	18	0.296	0.444
Finland	-0.684**	0.275	19	0.607	0.288
Greece	-0.507**	0.189	19	0.343	0.177
Iceland	-0.588**	0.215	15	0.439	0.222
Ireland	-1.280***	0.117	19	0.795	–
Portugal	-0.435**	0.180	19	0.306	0.143
Spain	-0.424***	0.106	19	0.626	0.138
Turkey	-0.379***	0.128	18	0.350	0.119
Australia	-0.242	0.166	19	0.110	0.069
New Zealand	-0.199	0.126	19	0.165	0.055
South Africa	-0.046	0.295	18	0.002	0.012
Bolivia	-0.368***	0.123	17	0.319	0.115
Chile	-0.303***	0.102	19	0.241	0.090
Colombia	-0.308*	0.148	19	0.178	0.092
Costa Rica	-0.441**	0.152	17	0.302	0.145
Ecuador	-0.259***	0.088	19	0.228	0.075
El Salvador	-0.265*	0.131	18	0.097	0.077
Honduras	-0.394*	0.216	17	0.144	0.125
Mexico	-0.577**	0.193	13	0.391	0.215
Uruguay	-0.270**	0.113	19	0.285	0.079
Venezuela, RB	-0.309	0.181	19	0.222	0.093
Trinidad and Tobago	-0.382	0.264	17	0.207	0.120
Iran, Islamic Rep.	-0.461*	0.234	19	0.158	0.155
Israel	-0.273	0.243	18	0.107	0.080
Jordan	-0.521**	0.204	18	0.284	0.184
Kuwait	-0.688***	0.173	17	0.514	0.291
Egypt, Arab Rep.	-0.328*	0.158	19	0.089	0.099
Sri Lanka	0.252	0.247	19	0.068	-0.056
India	-0.326*	0.186	19	0.117	0.099
Indonesia	-0.615***	0.162	16	0.553	0.239
Korea, Rep.	-0.801***	0.135	19	0.628	0.404
Malaysia	-0.708***	0.192	19	0.308	0.308
Pakistan	-0.379**	0.147	8	0.265	0.119
Philippines	-0.582**	0.217	19	0.291	0.218
Thailand	-1.151*	0.579	14	0.382	–
Ghana	-0.041	0.203	18	0.002	0.010
Kenya	-0.173	0.188	17	0.035	0.048
Mauritius	-0.108	0.246	15	0.010	0.028
Tanzania	-0.612**	0.227	12	0.419	0.237
Fiji	-0.269*	0.150	15	0.091	0.078

Notes: Robust standard errors clustered in parentheses; ***: significant at 1%; **: significant at 5%; *: significant at 10%. This table reports the results of regressing the growth of estimated technology parameter $\left(T_n^j\right)^{1/\theta}$ over the period from the 1960s to the 2000s on its initial value, by country. The speed of convergence, per decade, is reported in the last column. Missing values are due to the convergence coefficient being larger than 1.

Table A5. Country-by-Country Estimates of Relative Convergence, 1980s to 2000s

Country	β	s.e.	Obs.	R^2	Speed of Convergence, by decade
United Kingdom	-0.836***	0.203	19	0.478	0.904
Austria	-0.617*	0.316	19	0.354	0.480
Belgium-Luxembourg	-0.841***	0.219	19	0.489	0.919
Denmark	-0.778***	0.188	19	0.516	0.754
France	-1.164***	0.222	19	0.493	–
Germany	-0.698***	0.172	19	0.451	0.598
Italy	-0.303	0.355	19	0.074	0.181
Netherlands	-0.465**	0.217	19	0.244	0.312
Norway	-0.856***	0.108	19	0.781	0.969
Sweden	-0.519***	0.114	18	0.514	0.366
Switzerland	-1.106***	0.177	13	0.687	–
Canada	-0.516*	0.280	19	0.138	0.363
Japan	0.156	0.300	19	0.012	-0.073
Finland	-0.419*	0.212	19	0.343	0.271
Greece	-0.432***	0.128	19	0.531	0.283
Iceland	-0.706**	0.287	13	0.534	0.613
Ireland	-0.797**	0.313	19	0.320	0.797
Portugal	-0.230**	0.081	19	0.160	0.131
Spain	-0.401*	0.200	19	0.390	0.257
Turkey	-0.079	0.078	19	0.023	0.041
Australia	-0.015	0.255	19	0.000	0.008
New Zealand	0.022	0.171	19	0.001	-0.011
South Africa	-0.120	0.176	18	0.030	0.064
Argentina	-0.017	0.087	19	0.001	0.008
Bolivia	0.008	0.079	19	0.001	-0.004
Brazil	-0.273	0.250	16	0.131	0.160
Chile	-0.222**	0.081	19	0.252	0.125
Colombia	0.019	0.115	19	0.003	-0.010
Costa Rica	-0.356**	0.129	17	0.243	0.220
Ecuador	-0.222	0.136	19	0.126	0.125
El Salvador	0.023	0.240	18	0.001	-0.011
Honduras	-0.275	0.174	19	0.095	0.161
Mexico	-0.395*	0.189	18	0.165	0.251
Peru	0.150	0.100	19	0.099	-0.070
Uruguay	-0.137*	0.072	19	0.203	0.073
Venezuela, RB	0.249	0.187	19	0.072	-0.111
Trinidad and Tobago	0.031	0.154	18	0.002	-0.015
Iran, Islamic Rep.	0.536*	0.295	19	0.153	-0.215
Israel	0.094	0.124	18	0.032	-0.045
Jordan	-0.056	0.173	19	0.006	0.029
Kuwait	-0.259	0.201	17	0.091	0.150
Saudi Arabia	0.020	0.414	18	0.000	-0.010
Egypt, Arab Rep.	0.389	0.241	19	0.133	-0.164
Bangladesh	-0.024	0.146	17	0.002	0.012
Sri Lanka	0.031	0.063	19	0.008	-0.015
Taiwan Province of China	-0.115	0.258	19	0.014	0.061
India	-0.059	0.212	19	0.005	0.030
Indonesia	-0.241*	0.124	19	0.166	0.138
Korea, Rep.	-0.533*	0.282	19	0.235	0.380
Malaysia	-0.118	0.231	19	0.012	0.063
Pakistan	-0.188	0.253	8	0.074	0.104
Philippines	-0.158	0.229	19	0.024	0.086
Thailand	0.161	0.268	15	0.022	-0.075
Ethiopia	-0.246*	0.136	17	0.183	0.141
Ghana	-0.200	0.139	18	0.075	0.112
Kenya	0.068	0.124	17	0.015	-0.033
Mauritius	-0.019	0.130	18	0.001	0.010
Senegal	0.086	0.160	17	0.013	-0.041
Tanzania	0.157	0.292	12	0.044	-0.073
Fiji	-0.124	0.157	16	0.027	0.066
China	-0.160	0.190	19	0.037	0.087

Notes: Robust standard errors clustered in parentheses; ***: significant at 1%; **: significant at 5%; *: significant at 10%. This table reports the results of regressing the growth of estimated technology parameter $\left(T_n^j\right)^{1/\theta}$ over the period from the 1980s to the 2000s on its initial value, by country. The speed of convergence, per decade, is reported in the last column. Missing values are due to the convergence coefficient being larger than 1.

Table A6. $\theta = 4$: Pooled Regression Results

	(1) 1960s to 2000s	(2) 1980s to 2000s	(3) 1960s to 1970s	(4) 1970s to 1980s	(5) 1980s to 1990s	(6) 1990s to 2000s
Dep. Var: Log Change in $T^{1/\theta}$						
			Panel A: All Countries			
$\text{Log}\left(\text{Initial}T^{1/\theta}\right)$	-0.665***	-0.265***	-0.278***	-0.176***	-0.230***	-0.141***
	(0.045)	(0.031)	(0.030)	(0.025)	(0.029)	(0.039)
NB:						
Speed of convergence, per decade	*0.273*	*0.154*	*0.326*	*0.194*	*0.261*	*0.152*
Observations	929	1,122	991	1,074	1,183	1,335
R^2	0.680	0.647	0.678	0.666	0.743	0.672
			Panel B: OECD			
$\text{Log}\left(\text{Initial}T^{1/\theta}\right)$	-0.730***	-0.406***	-0.267***	-0.155***	-0.247***	-0.176**
	(0.094)	(0.071)	(0.043)	(0.031)	(0.046)	(0.072)
NB:						
Speed of convergence, per decade	*0.327*	*0.260*	*0.311*	*0.168*	*0.284*	*0.194*
Observations	393	405	396	394	407	410
R^2	0.755	0.709	0.785	0.662	0.627	0.673
			Panel C: non-OECD			
$\text{Log}\left(\text{Initial}T^{1/\theta}\right)$	-0.739***	-0.295***	-0.398***	-0.220***	-0.285***	-0.173***
	(0.054)	(0.044)	(0.041)	(0.034)	(0.040)	(0.053)
NB:						
Speed of convergence, per decade	*0.336*	*0.175*	*0.507*	*0.248*	*0.335*	*0.190*
Observations	536	717	595	680	776	925
R^2	0.733	0.632	0.708	0.687	0.754	0.683
Country FE	yes	yes	yes	yes	yes	yes
Sector FE	yes	yes	yes	yes	yes	yes

Notes: Standard errors clustered at the country level in parentheses; ***: significant at 1%; **: significant at 5%. This table reports the results of regressing the growth of estimated technology parameter $\left(T_n^j\right)^{1/\theta}$ on its initial value over different time periods and subsamples. The values of $\left(T_n^j\right)^{1/\theta}$ are estimated under the assumption that $\theta = 4$. The speed of convergence, per decade, is reported (in italics) underneath each coefficient estimate.

Table A7. $\theta = 4$: Welfare Gains in the Single-Country Counterfactual Relative to Baseline

	Median	St. Dev.	Min	Max	Countries
OECD	0.014	0.025	-0.018	0.086	22
Non-OECD	0.045	0.106	-0.098	0.602	53
NB: Overall gains from trade					
OECD	0.112	0.072	0.020	0.266	
Non-OECD	0.086	0.057	0.010	0.258	

Notes: This table reports the proportional change in welfare under the counterfactual scenario with respect to the baseline, under the assumption that $\theta = 4$. The counterfactual assumes that for each individual country, comparative advantage remained as it was in the 1960s, while its T's grew at the same country-specific average rate between the 1960s and the 2000s. All other countries' comparative advantage is taken from the data. In the baseline comparative advantage is as it is in the data for the 2000s. The lower panel reports the total gains from trade relative to autarky in the baseline for the 2000s.

Table A8. $\theta = 4$: Welfare Gains in the Global Counterfactual Relative to Baseline

	Median	St. Dev.	Min	Max	Countries
Panel A: CA fixed to 1960s in all countries					
OECD	0.015	0.022	-0.013	0.077	22
Non-OECD	0.031	0.102	-0.129	0.569	53
Panel B: CA fixed to 1960s in OECD countries only					
OECD	0.014	0.024	-0.018	0.080	
Non-OECD	0.000	0.003	-0.010	0.007	
Panel C: CA fixed to 1960s in non-OECD countries only					
OECD	0.001	0.003	-0.002	0.010	
Non-OECD	0.034	0.104	-0.129	0.589	

Notes: This table reports the proportional change in welfare under the counterfactual scenario with respect to the baseline, under the assumption that $\theta = 4$. The counterfactual assumes that in all countries in the world (Panel A), in OECD (Panel B) and the non-OECD (Panel C), comparative advantage remained as it was in the 1960s, while its T's grew at the same country-specific average rate between the 1960s and the 2000s. In the baseline comparative advantage is as it is in the data for the 2000s.

Table A9. Comparison to Measured TFP from STAN Database

Panel A: Sector-by-Sector Rank Correlations

ISIC code	Sector Name	Correlation	Countries
15	Food and Beverages	0.6000	5
16	Tobacco Products	0.6000	5
17	Textiles	0.6429	7
18	Wearing Apparel, Fur	0.4286	6
19	Leather, Leather Products, Footwear	0.2000	6
20	Wood Products (Excl. Furniture)	0.4818	11
21	Paper and Paper Products	0.2500	9
22	Printing and Publishing	0.6833	9
23	Coke, Refined Petroleum Products, Nuclear Fuel	0.6667	8
24	Chemical and Chemical Products	0.4667	10
25	Rubber and Plastics Products	0.9364	11
26	Non-Metallic Mineral Products	0.8091	11
27	Basic Metals	0.1833	9
28	Fabricated Metal Products	0.8833	9
29C	Office, Accounting, Computing, and Other Machinery	0.7343	12
31A	Electrical Machinery, Communication Equipment	0.5952	8
33	Medical, Precision, and Optical Instruments	0.9048	8
34A	Transport Equipment	0.6000	9
36	Furniture and Other Manufacturing	0.7143	7

Panel B: Fixed Effects Regression

	(1)	(2)	(3)
Dep. Var: Log Sectoral Productivity Implied by Sectoral Measured TFP			
$\log\left(T_n^j\right)^{1/\theta}$	1.188***	1.457***	0.751*
	(0.222)	(0.145)	(0.397)
Observations	160	160	160
R-squared	0.116	0.689	0.849
Partial ρ	0.341	0.583	0.175
Sector FE	no	yes	yes
Country FE	no	no	yes

Notes: This table reports the results of comparing the productivity estimates using the main procedure adopted in the paper ($\left(T_n^j\right)^{1/\theta}$) with TFP estimated directly using production data from the OECD STAN database. Panel A reports the Spearman rank correlations of the two alternative productivity measures by sector. Panel B reports the results of a fixed effects regression of directly measured TFP from STAN on $\left(T_n^j\right)^{1/\theta}$. In Panel B, robust standard errors in parentheses; *: significant at 10%; ***: significant at 1%. "Partial ρ" is the partial correlation between the right-hand side and the left-hand side variables, after netting out the fixed effects included in the column.

Figure A1. Comparison of Aggregate Measured TFP and Weighted Average $T^{1/\theta}$

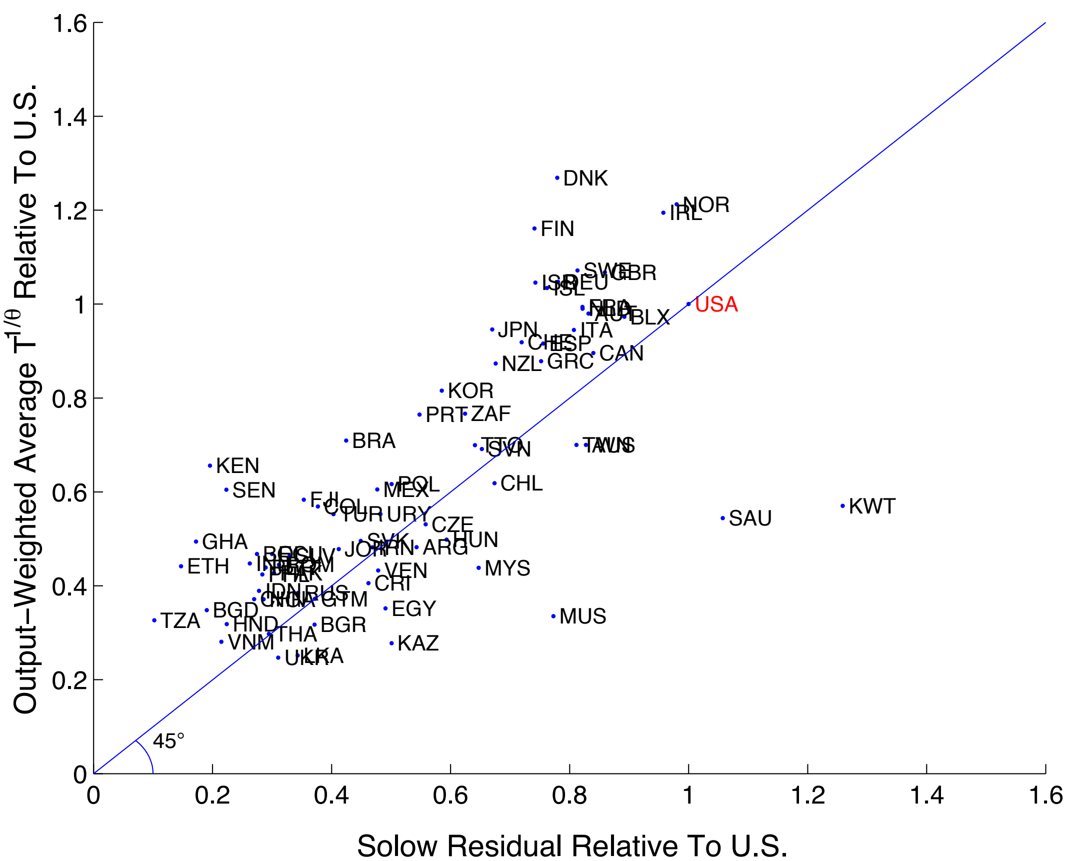

Notes: This figure presents the scatterplot of measured TFP computed using standard methods based on the Penn World Tables on the x-axis against the output-weighted average of $T^{1/\theta}$ on the y-axis, for the 2000s. The solid line is the 45-degree line.