



Australian School of Business

Working Paper

Never Stand Still

Australian School of Business

Australian School of Business Research Paper No. 2014 ECON 26

Innovation, Product-Cycle Trade, and the Cross-Country Distribution of Income

Scott French

This paper can be downloaded without charge from
The Social Science Research Network Electronic Paper Collection:

<http://ssrn.com/abstract=2432806>

Innovation, Product-Cycle Trade, and the Cross-Country Distribution of Income

Scott French*

April, 2014

Abstract

This paper develops a quantitative, multi-country model of endogenous growth, international trade, and international knowledge flows in order to understand how access to both foreign products and technologies, together, influences innovation incentives and the world distribution of income. An endogenous product cycle arises in equilibrium, in which innovative countries engage in both horizontal and vertical research, while others far from the technological frontier specialize in learning about and applying research previously conducted abroad. The effect of trade barriers on the level and dispersion of income across countries is found to be larger than would be predicted by a static trade model, and the effect of access to international knowledge flows is also quantitatively important and dependent on trade flows. For instance, halving the cost of learning reduces income dispersion by 23%, while doing so after eliminating asymmetric international trade barriers reduces income dispersion by only 10%.

JEL Classification: F11, F12, F14, O19, O31, O33, O40

Keywords: Income differences, Trade, Endogenous growth, Product cycle, Innovation, Productivity, Technology diffusion

*School of Economics, University of New South Wales, scott.french@unsw.edu.au. I am thankful to Dean Corbae for his guidance and support. I am also thankful to Arghya Ghosh, Tess Stafford, Kripa Freitas, Kim Ruhl, and Jonathan Eaton for helpful comments. All errors are my own. A previous version of this paper was titled “Innovation and Product-Level Trade”

1 Introduction

It is well known that differences in income per worker across countries are very large and that, for the majority of countries, these differences have been persistent over time. A great deal of work has been devoted to determining the sources of these differences. Accounting exercises based on the Solow-Swan neoclassical growth model, such as [Caselli \(2005\)](#), find that productivity differences account for the majority of income differences. While such studies tend to treat countries in isolation and take productivity as an exogenous factor, others have sought to measure the contribution of endogenous innovation (e.g. [Howitt, 2000](#)), technology diffusion (e.g. [Eaton and Kortum, 1999](#)), and international trade (e.g. [Vaugh, 2010](#)) to measured productivity.

In this paper, I assess the contribution of all three of these together in a unified framework in which each is determined endogenously. To this end, I develop a quantitative, multi-country model of endogenous growth, international trade, and international knowledge flows and use this model to quantify the effect of countries' access to foreign products and technologies on the worldwide distribution of income. An endogenous product cycle emerges in equilibrium, in which most products, and techniques for producing them, are invented in the richest countries. Poorer countries benefit from producing and exporting older products using knowledge that has been developed in the richer countries. I fit the model to data on international trade flows and research expenditure and argue that the ease of access to foreign products and knowledge, and its effects on innovation incentives, is quantitatively important in accounting for the cross-country differences in productivity and income that exist in the data. For example, the elimination of asymmetric trade barriers causes a 68% decrease in the dispersion of income per worker across countries, and the change is 20% larger than would be predicted by a model that did not account for the effect on trade barriers on countries' research expenditure.¹

Innovation in the model, as in [Young \(1998\)](#), takes the form of both horizontal and vertical research and development – i.e. inventing new products and developing new techniques for producing existing products, respectively.² In the baseline model, similar to [Rivera-Batiz and Romer \(1991\)](#), knowledge of the former type – that a new product has been invented – travels instantly and costlessly throughout the world, but in order to be able to produce an existing product, a potential entrant must develop a production technique either through original research (vertical R&D) or by devoting resources to learning and applying knowledge previously created through the vertical R&D of others.³ International trade is driven by Ricardian forces – buyers in each country

¹In the specification used in this paper, trade barriers are asymmetric due to border costs that vary by exporter. As [Vaugh \(2010\)](#) argues, such barriers, unlike typical variables from the gravity literature, which tend to be symmetric, are likely to be policy-related.

²[Howitt \(1999\)](#), [Peretto \(1998\)](#), and [Dinopoulos and Thompson \(1998\)](#) also develop endogenous growth models with both horizontal and vertical innovation. While these models were primarily concerned with scale effects on growth, the model of this paper features both forms of innovation to highlight the differential effects of international trade and knowledge flows on the returns to each.

³This assumption is in the spirit of [Nelson \(1982\)](#), who argues that there is information contained in previously conducted research – the “logy” component of technology – that is public in nature and can be utilized by firms to develop production techniques without conducting original research.

purchase a given product from the producer that can produce and deliver it at the lowest cost – and the productivity with which a product is produced in a given country is a function of the history of innovative activity there since the the time of the product’s invention.

The model admits a balanced growth path in which every country grows at a constant rate, and the share of a country’s output devoted to each type of innovative activity, along with its average research productivity, its size, and the trade barriers it faces, determines its level of income relative to other countries.⁴ Along the balanced growth path, the expected productivity with which a product can be produced can be characterized as a function of the age of the product and the location of its inventor. Initially, only the inventor can produce the product, but over time better techniques for producing it are developed, and the location of the low-cost provider evolves as new vertical R&D related to it accumulates and is learned and applied around the world.

There are two major forces that lead to the emergence of a product cycle in equilibrium. First, because knowledge created through vertical research is only useful in countries where it has not been previously applied, the returns to learning, as in [Acemoglu et al. \(2006\)](#), are increasing in a country’s average distance to the technological frontier, which is decreasing in a country’s average research productivity and the size of its labor force relative to the rest of the world. Second, with international trade, expected profits from producing an older product, for which vertical R&D has accumulated and been applied around the world, are relatively lower for producers in high-wage countries compared with newly invented products for which these producers face no direct competition initially. Because relative wages are also increasing in research productivity in equilibrium, this implies that expenditure on horizontal R&D is relatively greater in high-income countries, expenditure on learning is relatively greater in low-income countries, and relative expenditure on vertical research depends on the relative strength of these two effects. Thus, products tend to be invented in high-income countries, which produce and export products early in their life cycle, while low-income countries have a comparative advantage in older products for which knowledge generated by foreign vertical research has been learned over time.

The model of innovation and international trade is built around the framework of [Eaton and Kortum \(2001\)](#), which combines the search-theoretic model of innovation and growth from [Kortum \(1997\)](#) with the multi-country Ricardian trade model of [Eaton and Kortum \(2002\)](#) into a tractable multi-country model of international trade and growth. To this, I introduce the structure of innovation and international knowledge flows described above. In this setting, unlike in [Eaton and Kortum \(2001\)](#), research expenditure varies across countries at different positions in the world income distribution and with differing trade barriers. I embed this model into a neo-classical growth framework, as [Alvarez and Lucas \(2007\)](#) do in a static context using the model of [Eaton and Kortum \(2002\)](#). As a result, the model generates predictions for both nominal and real GDP, productivity, and international trade flows that are directly comparable with the data and development accounting exercises based on the neoclassical growth model.

⁴The model imposes the restriction that *relative* productivity across each innovative activity – horizontal research, vertical research, and learning – is the same across countries, so specialization in these activities is not driven by such exogenous differences.

To quantify the effects of international trade and knowledge flows on the cross-country distribution of income, I use data on manufacturing output and bilateral trade flows to estimate bilateral trade barriers and levels of average research productivity across countries, and I use data on research and development expenditure to infer the relative productivity of resources across each type of innovative activity. The parameter values recovered imply that developing a new production technique via original research requires a 9% smaller investment, on average, than bringing a new product to market, while applying foreign research requires an investment 60% smaller than conducting original vertical research. I then show that the model is able to accurately predict differences in real income per worker across countries, the shape of the relationship between R&D expenditures and income per worker in the data, and salient features of the data on product-level trade flows.

Next, I perform a series of counterfactual experiments designed to demonstrate the effects of international trade and knowledge flows on the cross-country distribution of income. In the first set of experiments, I vary the level of trade barriers, considering both a move to a world with no trade flows and the elimination of all asymmetric trade barriers. When only the static effect of these changes is considered, the results are similar to [Vaugh \(2010\)](#). The elimination of asymmetric trade barriers, which are estimated to be systematically greater for poorer countries, reduces income dispersion – measured as the 90/10 ratio of income per worker – by more than half, while the move to trade autarky has a relatively small effect on cross-country income differences. When the effect on research incentives is taken into account, the effect is amplified. The reduction in trade barriers causes the most innovative countries to devote more resources to horizontal research, which benefits all countries, and the static gains of poorer countries allow them to increase gross expenditure on all innovative activities. As a result, income differences fall by a further 20%, and total world income increases by an additional 80%, so that world income more than doubles, overall, in response to the reduction in trade barriers.

The second set of experiments vary the ease of international knowledge flows. The elimination of access to knowledge created by foreign vertical research greatly increases expenditure on both horizontal and vertical innovation, implying a substantial increase in both the number of products available and the technological frontier for each product, but average productivity falls everywhere because resources devoted to vertical research produce fewer techniques than if they had been devoted to learning, especially for countries far from the technological frontier. The result is that world income falls by half, and income differences increase by 18%. A halving of the cost of learning causes changes of a similar magnitude in the opposite direction. International trade and knowledge flows are found to be substitutes in their effects on the distribution of world income. The increase in total world income from introducing learning is slightly higher in a world without trade than with the baseline trade costs, and the introduction of learning causes a 29% decrease in the 90/10 ratio of income per worker in a world without trade and only a 16% decrease in a world with the baseline level of trade costs. Finally, the elimination of access to knowledge of the existence of products invented abroad is devastating to countries' incomes, especially for countries that are

small and relatively unproductive at research. Compared to a setting in which there are already no trade or vertical knowledge flows, moving to complete isolation entails a further 95% reduction of world income and a nearly eight-fold increase in income differences, with much greater effects for the worst affected. While these predictions may seem extreme they are in line with historical evidence on small, isolated communities, for example the case of Tasmania in [Diamond \(1997\)](#).

While this paper is primarily focused on measuring the contribution of international trade and knowledge flows to the world income distribution, it also contributes to a large literature going back to [Vernon \(1966\)](#) on product cycles and technology transfer in international trade and foreign direct investment. This paper is most related to the strand of literature, going back to [Krugman \(1979\)](#), devoted to modelling the causes and consequences of product cycles in international trade, particularly [Grossman and Helpman \(1991a,b\)](#), which develop theories of endogenous product cycles. However, to my knowledge, this is the first paper to develop a quantitative, multi-country model in which product cycles arise entirely endogenously and not because certain countries are assumed to have an ex-ante advantage in either innovation or imitation. This paper has less in common with [Antràs \(2005\)](#), which explicitly models the international frictions that prevent a given product from being produced with the best technique anywhere in the world. Instead, this paper is largely complementary in specifying such frictions in a simplified manner – assuming useful vertical knowledge can only be applied through costly learning by local firms – in order to develop a model much more amenable to making quantitative predictions about trade flows and the world income distribution in a multi-country setting.

The next section develops the model, first detailing the structure of the world economy before explaining how research expenditures over time determine the distribution of productivity across products and countries and how this, in turn, determines prices and bilateral trade flows. Section 3 characterizes the general equilibrium, first describing the equilibrium at a point in time given the history of research expenditures and then characterizing the balanced growth path along which these expenditures are determined. Section 4 describes the choice of parameters, properties of the numerically computed equilibrium, and ability of the model to predict salient features of the data. Section 5 presents the results of the counterfactual experiments, and the final section concludes.

2 Model

Time is continuous. The world economy is made up of N countries. Each consists of a continuum of workers with mass $L_i(t)$ at time t , which grows exogenously and exponentially at the common rate g_L . Each worker is endowed with a single unit of labor which is supplied inelastically.

2.1 Preferences and Final Good Production

A representative household in each country maximizes the expected, discounted future flow of utility given by

$$U_i = \int_0^{\infty} e^{-(r-g_L)t} c_i(t) dt, \quad (1)$$

where $c_i(t)$ is per-capita consumption of the final good, which is non-traded. The final good is produced using labor and a composite manufactured good according to

$$Y_i(t) = L_i^f(t)^\gamma Q_i^f(t)^{1-\gamma}, \quad (2)$$

where $L_i^f(t)$ and $Q_i^f(t)$ are the respective measures of labor and the composite good devoted to final good production. In what follows, the time variable is suppressed where there is no ambiguity.

2.2 Intermediate Good Production

The composite manufactured good is produced by combining manufactured intermediate goods according to a symmetric Spence-Dixit-Stiglitz aggregator. At a point in time, there is a continuum of such intermediates available, labelled $j \in [0, J(t)]$. Each intermediate good can potentially be produced by a number of firms in any country, but such varieties of a good are assumed to be perfect substitutes so that the aggregate output of the composite intermediate good is given by

$$Q_i = \left(\int_0^{J(t)} (q_i^j)^{\frac{\sigma-1}{\sigma}} dj \right)^{\frac{\sigma}{\sigma-1}}, \quad (3)$$

where q_i^j is the quantity of intermediate j (from all sources) that is used, and $\sigma > 1$ is the elasticity of substitution across intermediates.

Individual intermediate goods are produced using labor and the composite intermediate good by imperfectly competitive firms that differ in terms of total factor productivity. It is convenient to rank firms according to their productivity, in which case the k th most productive firm at producing intermediate product j in country i has the following Cobb-Douglas production function:

$$x_i^{jk} = Z_i^{jk} (L_i^{jk})^\beta (Q_i^{jk})^{1-\beta}. \quad (4)$$

2.3 Market Structure

The technologies for producing the final good and for combining individual intermediate goods into the composite good are assumed to be common knowledge, and there are no barriers to entry. Therefore, both sectors can be characterized by a cost-minimizing representative firm. Thus, demand for intermediate good j (in value terms) in country n is

$$X_n^j \equiv p_n^j q_n^j = \left(\frac{p_n^j}{P_n} \right)^{1-\sigma} X_n, \quad (5)$$

where p_n^j is the lowest price in n of intermediate j , $X_n = P_n Q_n$ is total expenditure on intermediate goods in n , and P_n is the price of the composite good in n , given by

$$P_n = \left(\int_0^{J(t)} (p_n^j)^{1-\sigma} \right)^{\frac{1}{1-\sigma}}. \quad (6)$$

Given the value of P_n , the price of the final good in n can be expressed as

$$P_n^y = \tilde{\gamma} w_n^\gamma P_n^{1-\gamma}, \quad (7)$$

where w_n is the wage in n , and $\tilde{\gamma} = \gamma^{-\gamma}(1-\gamma)^{\gamma-1}$.⁵

Producers of individual intermediate goods enjoy proprietary access to their technologies and engage in Bertrand competition. As a result, only the producer which can deliver its product at the lowest cost will operate in equilibrium, and, to prevent entry by other firms, will charge a price no higher than the second-lowest cost of delivering the product. Cost minimization by firms implies that the k th lowest cost way of producing j in i is

$$C_i^{jk} = \frac{\tilde{\beta} w_i^\beta P_i^{1-\beta}}{Z_i^{jk}} \equiv \frac{c_i}{Z_i^{jk}} \quad (8)$$

where $\tilde{\beta} = \beta^{-\beta}(1-\beta)^{\beta-1}$. This market structure implies that, at most, the costs of the two lowest-cost firms are relevant in equilibrium.

2.4 International Trade

While the final and composite manufactured goods are assumed to be non-tradable, intermediate goods can be shipped anywhere in the world subject to trade barriers that take the convenient “iceberg” form, as in [Samuelson \(1954\)](#), meaning that delivering 1 unit of a good from i to n requires shipping $d_{ni} > 1$ units. Combined with (8), this implies that the k th lowest cost way of delivering a unit of j from i to n is

$$C_{ni}^{jk} = \frac{c_i d_{ni}}{Z_i^{jk}}. \quad (9)$$

In equilibrium, country n will be served by the firm with the lowest cost of delivering j ,

$$C_n^{j1} = \min_i \{C_{ni}^{j1}\},$$

which will charge a price that is the lesser of the second lowest cost of delivering j and the profit-maximizing monopoly mark-up over its marginal costs:

$$p_n^j = \min\{C_n^{j2}, \bar{m} C_n^{j1}\}, \quad (10)$$

⁵In what follows, without loss of generality, I normalize w_1 to 1 at every point in time. However, for both clarity and convenience, I leave the wage term in all expressions for all countries.

where $\bar{m} = \sigma/(\sigma - 1)$, and $C_n^{j2} = \min\{C_{ni^*}^{j2}, \min_{i \neq i^*}\{C_{ni}^{j1}\}\}$, where $i^* = \operatorname{argmin}_i\{C_{ni}^{j1}\}$.

2.5 Research and Development

There is a continuum of potential entrants into the intermediates goods industry. To enter, a firm must successfully invent a new product or develop a technique for producing an existing product. To do this, a firm engages in one of three innovative activities, which I term *horizontal* research, *vertical* research, and *learning*. Horizontal research involves designing a new type of intermediate input and, when successful, results in an expansion of the measure of available products, $J(t)$. Vertical research involves searching for new techniques for producing existing intermediates, which may be more productive than existing techniques. I assume that the existence of a new product is immediately known throughout the world but that firms are able to prevent competitors in any country from exactly copying their production techniques. Potential entrants can, however, learn about the vertical research itself and utilize this information in an attempt to develop a production technique. This is the third type of innovative activity. The key distinction between vertical research and learning is that vertical research generates new knowledge that is useful in developing production techniques, while learning only applies knowledge previously created.

Potential entrants may invest one unit of final output into one of the three types of research project. I assume that potential entrants differ in terms of the likelihood of success of such an investment and choose whether to invest sequentially, beginning with the most productive entrant for each type of activity. If a measure M_i^x of firms in country i invest in innovative activity of type $x \in h, v, \ell$ – which correspond to horizontal research, vertical research, and learning, respectively – the probability of success of the marginal firm’s project is

$$\lambda \alpha_i^x (M_i^x)^{\lambda-1}, \quad (11)$$

where $\lambda < 1$.⁶ This implies that the aggregate flow of successful projects of type x in i at time t is

$$R_i^x(t) = \alpha_i^x (s_i^x(t) Y_i(t))^\lambda, \quad (12)$$

where $s_i^x(t)$ is the share of final output devoted to the an innovative activity.

I further assume that $\alpha_i^x = \alpha_i \alpha^x$. Under this assumption, α^x reflects the relative productivity of a unit of output in each innovative activity – which is plausibly constant across countries – and α_i reflects any country-specific characteristics that influence the returns to all types of innovative activity, such as the endowment of human capital, as in [Benhabib and Spiegel \(1994\)](#), or policies that promote or create barriers to innovation or learning about foreign technology, as in [Parente and Prescott \(2002\)](#). Maintaining this assumption has the advantage of both greatly reducing the number of parameters that must be chosen in the quantitative analysis that follows and introducing

⁶This specification arises if the distribution of productivity of potential entrants for each type of research is Pareto with shape parameter $1/(1 - \lambda)$ and lower bound proportional to $\alpha_i^x (M_i^x)^{1-\lambda}$. Thus, a greater value of α_i^x shifts the entire productivity distribution to the right, while an increase in M_i^x shifts the distribution to the left by appending less productive entrants.

a degree of discipline to the model by ensuring that specialization across countries into different types of innovation is not driven by exogenous differences in relative productivity across these activities but must arise endogenously in general equilibrium.

In equilibrium, free entry will drive the expected payoff of conducting R&D to zero, which implies the following conditions:

$$\lambda \alpha_i^r (s_i^x(t) Y_i(t))^{\lambda-1} V_i^x(t) = P_i^y(t), \quad \forall i, x, \quad (13)$$

where $V_i^x(t)$ is the present value of expected future profits resulting from a project of type x , to be determined below.

2.5.1 Horizontal Innovation

The aggregate rate at which new intermediate inputs are invented by researchers in i is simply the flow of successful type h projects, thus

$$\dot{J}_i(t) = R_i^h(t).$$

This implies that the total measure of intermediates available in the world grows at rate

$$\dot{J}(t) = \sum_i R_i^h(t). \quad (14)$$

2.5.2 Vertical Innovation

I assume that vertical research is undirected, so a new technique is equally likely to apply to any existing product. Thus, the aggregate arrival rate of new production techniques for a given product is

$$\dot{T}_i^{vj}(t) = \frac{R_i^v(t)}{J(t)}. \quad (15)$$

I also assume that the inventor of a new product receives an expected value of one technique for producing it at the time of invention, which implies that the expected stock of techniques available in i for producing j at the time of its invention, \underline{t}^j , is

$$T_i^{vj}(\underline{t}^j) = \mathbf{1}_{im^j}, \quad (16)$$

where m^j is the country in which j was invented, and $\mathbf{1}_{im^j}$ is the indicator function which is equal to 1 if i is equal to m^j .⁷

⁷Allowing the number of techniques that arrive with a new invention to be uncertain greatly simplifies the analysis that follows. One can think of the process of invention as involving simultaneous horizontal and vertical research. The vertical research may yield more than one useable technique before the design of the product is entirely complete, or knowledge of the product may leak into the public sphere before a production technique can be implemented.

2.5.3 Learning

I assume that learning is also undirected, which implies that the probability that a potential entrant learns about research pertaining to a particular product is proportional to the share of the total world stock of accumulated vertical research which pertains to that product. I further assume that only research that has yet to be applied in a given country is useable.⁸ Thus, the aggregate arrival rate of techniques for a given product is given by

$$\dot{T}_i^{\ell j}(t) = R_i^\ell(t) \frac{T_w^j(t) - T_i^{vj}(t)}{T_w(t)}, \quad (17)$$

where $T_w^j(t) = \sum_i T_i^{vj}(t)$ is the worldwide stock of original research into techniques for producing product j , and $T_w(t) = \int_0^{J(t)} T_w^j(t) dj$ is the total world stock of original research applying to all products.

Combining (15) and (17), the arrival rate of techniques for a given product via any activity by researchers in i is

$$\dot{T}_i^j(t) = \dot{T}_i^{vj}(t) + \dot{T}_i^{\ell j}(t). \quad (18)$$

This law of motion for the number of techniques available in a given country resembles that for productivity in [Vandenbussche et al. \(2006\)](#) in that the local technology frontier increases both through domestic innovation and through learning – in their terminology, adoption or imitation – where the strength of the latter is increasing in a country’s distance to the world technological frontier.⁹

Equation (18) together with the initial condition for T_i^{vj} , equation (16), implies that, at a given point in time, the number of techniques available in i for producing j is a Poisson random variable with parameter

$$T_i^j(t) = \mathbf{1}_{im^j} + \int_{t^j}^t \dot{T}_i^j(s) ds. \quad (19)$$

2.6 Productivity

The productivity associated with a particular production process for a given product is the realization of a Pareto random variable, with shape parameter $\theta > 0$.¹⁰ I show in [Appendix D](#) that this implies that the joint distribution of the highest and second-highest levels of productivity for producing product j in i is given by

$$F_i^j(z_1, z_2) = Pr\left(Z_i^{j1} < z_1, Z_i^{j2} < z_2\right) = \left[1 + T_i^j(z_2^{-\theta} - z_1^{-\theta})\right] \exp\left\{-T_i^j z_2^{-\theta}\right\}. \quad (20)$$

⁸Under the assumption of Bertrand competition, this is equivalent to assuming that a common piece of research leads to production techniques with the same level of productivity for producers in the same country but could lead to a technique with a different level of productivity when applied in another country. [French \(2013\)](#) considers the more complicated case in which the productivity of techniques developed from the same original research is imperfectly correlated both within and across countries.

⁹The law of motion for productivity in [Vandenbussche et al. \(2006\)](#) is based on similar specifications explored in [Benhabib and Spiegel \(1994\)](#) and [Acemoglu et al. \(2006\)](#).

¹⁰That is, $\Pr(Z > z) = (z/\underline{z})^{-\theta}$ for $z \in (\underline{z}, \infty)$.

This is a product-level analogue of the distribution detailed in [Bernard et al. \(2003\)](#), which is a two-variable generalization of the Fréchet (type-II extreme value) distribution used in [Eaton and Kortum \(2002\)](#).¹¹ The parameter T_i^j governs the expected level of productivity with which j is produced in i , while θ governs the degree of dispersion in product-level productivity, with higher values of θ implying lower variance. In other words, the average level of T_i^j across products determines i 's overall absolute advantage, variation in T_i^j across products governs i 's deterministic pattern of comparative advantage due to its history of research, and θ determines the degree of comparative advantage across countries due to idiosyncratic productivity differences.

2.7 Prices and Trade Flows

Given expressions for demand for intermediate inputs (5), the costs of delivering a given product to a given destination (9), the price that will be charged for this product (10), and the distribution of productivity across products and source countries (20), the following results can be derived, which will allow the complete characterization of prices and bilateral trade flows.

1. The joint distribution of the lowest and second-lowest cost of delivering j to n from anywhere in the world is characterized by the following complementary cdf:

$$\tilde{G}_n^j(c_1, c_2) = \Pr(C_n^{j1} > c_1, C_n^{j2} > c_2) = \left[1 + \Phi_n^j(c_2^\theta - c_1^\theta)\right] \exp\left\{-\Phi_n^j c_2^\theta\right\},$$

where $\Phi_n^j = \sum_i T_i^j (c_i d_{ni})^{-\theta}$.

2. The probability that country i is the lowest-cost source of product j for country n is

$$\pi_{ni}^j = \frac{T_i^j (c_i d_{ni})^{-\theta}}{\Phi_n^j}. \quad (21)$$

3. The joint distribution of the lowest and second-lowest cost of delivering j to n , *conditional on the low-cost seller being from i* , is also given by $\tilde{G}_n^j(c_1, c_2)$. Thus, the costs of products that i actually sells in n are identically distributed for all countries, with lower-cost sources selling more products on average.
4. The distribution of the ratio C_n^{j2}/C_n^{j1} is

$$\Pr(C_n^{j2}/C_n^{j1} < m) = 1 - m^{-\theta},$$

and the distribution of the markup in any destination is independent of the identity of the product or the location of the low-cost producer.

¹¹The marginal distribution of z_1 , obtained by replacing z_2 with z_1 in (20) is $F_i^j(z_1) = \exp\{-T_i^j z_1^{-\theta}\}$, which is the product-level analogue of equation (4) in [Eaton and Kortum \(2002\)](#).

5. The price index for the composite intermediate good in n , defined by (6), is

$$P_n = \hat{\gamma} \Phi_n^{-\frac{1}{\theta}}, \quad (22)$$

where

$$\Phi_n = \left(\int_0^J (\Phi_n^j)^{\frac{\sigma-1}{\theta}} dj \right)^{\frac{\theta}{\sigma-1}},$$

given that $\theta > \sigma - 1$.¹² The parameter $\hat{\gamma}$ is a collection of constants.¹³

6. The expected value of nominal expenditure by n on product j , given by (5), relative to total expenditure on intermediates is

$$\frac{E[X_n^j]}{X_n} = \left(\frac{\Phi_n^j}{\Phi_n} \right)^{\frac{\sigma-1}{\theta}}. \quad (23)$$

7. Expected profit of a firm from selling in n , conditional on being the low-cost provider of a given product, is a constant share, $1/(1+\theta)$, of expected sales. This share is thus independent of the identity of the destination, the identity of the product, and the location of the low-cost provider.

Proofs of these results are given in Appendix D. An immediate implication of results 2, 3, and 6 is that the expected value of sales of product j to destination n from source i is given by

$$E[X_{ni}^j] = \frac{T_i^j (c_i d_{ni})^{-\theta}}{\Phi_n^j} \left(\frac{\Phi_n^j}{\Phi_n} \right)^{\frac{\sigma-1}{\theta}} X_n. \quad (24)$$

Understanding the role of Φ_n^j in this expression provides a great deal of intuition about the effect of research activity on trade flows and income. This parameter summarizes the amount of research and learning that pertains to j that has accumulated worldwide, weighted by the costs of producing a good in each source and delivering it to n . From equation (21), we can see that, given the costs of production and stock of research pertaining to j in i , a greater value of Φ_n^j implies a lower probability that the low-cost provider of j to n is in i . From (23), we can see that Φ_n^j also plays another role. A higher value of Φ_n^j implies that, in expectation, a greater share of expenditure in n is devoted to j . This is because a higher value of Φ_n^j implies that the price of j

¹²This assumption is required for P_n to be well defined. This is because, if intermediate inputs are sufficiently substitutable (high σ), and the dispersion of productivity is sufficiently great (small θ), then producers of the composite intermediate will concentrate expenditure on a small set of products with prices arbitrarily close to zero.

¹³Specifically, $\hat{\gamma} = \left[\Gamma \left(\frac{2\theta+1-\sigma}{\theta} \right) \left(1 + \frac{\sigma-1}{\theta-(\sigma-1)} \bar{m} \right) \right]^{\frac{1}{1-\sigma}}$, where $\Gamma(\cdot)$ is the gamma function. Note that γ is used elsewhere in this paper for other purposes. Also, note that, because Φ_n^j is a function of the full set of intermediate price indexes, these price indexes are only implicitly defined by (22). Alvarez and Lucas (2007) show that this system has a unique solution.

in n is expected to be lower – with θ governing the elasticity of this effect because a higher value of θ implies less variation in productivity draws, meaning that it takes more draws, on average, to produce a given fall in the realized lowest cost. The lower expected price then translates to higher expected expenditure with elasticity $\sigma - 1$. Thus, there are two opposing effects of Φ_n^j on expected sales.

Since it must be the case that $\theta > \sigma - 1$ in order for P_n to be well-defined, the former effect dominates. Thus, sales by firms in i of a particular product will tend to be greater the lower is the stock of research pertaining to the product that has accumulated in other countries as well as the greater the degree to which the research stock is concentrated in countries with high costs of production and transportation.

2.7.1 Aggregate Trade Flows

By integrating (24) over j , we arrive at an expression for the share of expenditure by n on products purchased from i :

$$\pi_{ni} \equiv \frac{X_{ni}}{X_n} = \tilde{T}_{ni} \frac{T_i (c_i d_{ni})^{-\theta}}{\Phi_n}, \quad (25)$$

where

$$T_i = \left(\int_0^{J(t)} \left(T_i^j \right)^{\frac{\sigma-1}{\theta}} dj \right)^{\frac{\theta}{\sigma-1}}, \quad (26)$$

and

$$\tilde{T}_{ni} = \int_0^{J(t)} \frac{T_i^j}{T_i} \left(\frac{\Phi_n^j}{\Phi_n} \right)^{\frac{(\sigma-1)}{\theta}-1} dj. \quad (27)$$

The model of [Eaton and Kortum \(2002\)](#) yields a similar expression to (25), except that the latter includes the term \tilde{T}_{ni} . While T_i is an index of the total number of techniques across all products that are available in i , \tilde{T}_{ni} summarizes the effect on trade flows of the interaction of the distribution of these techniques in i across products with the distributions of the other countries of the world. If techniques were distributed uniformly across all products, then \tilde{T}_{ni} would disappear from (25). In general, however, (27) shows that a country will tend to sell more if its products in a given destination if its production techniques tend to be concentrated in products for which Φ_n^j is smaller. In other words, a country will tend to export more if its stock of research tends to pertain to a relatively unique set of products.

2.7.2 Equilibrium Wages

Expression (25) can be manipulated to give a useful expression for the real value of wages (in terms of the composite intermediate good):

$$\frac{w_i}{P_i} = \left(\frac{T_i}{\pi_{ii}} \right)^{\frac{1}{\beta\theta}} \tilde{T}_{ii}^{\frac{1}{\beta\theta}}. \quad (28)$$

While this is not a solution to the general equilibrium, as π_{ii} is an endogenous variable, it is informative as to how countries' histories of research determine the world distribution of income as well as the gains from trade. The first term is typical of trade models based on [Eaton and Kortum \(2002\)](#). It demonstrates that real wages are greater for countries that are more productive on average – i.e. have a greater absolute advantage – and for countries that trade a greater share of their production, which allows them to specialize in a narrower set of products for which they are relatively more productive.

However, there are two key differences between the way in which technology and trade barriers influence the world income distribution in this model compared with a standard Eaton-Kortum framework. The first can be seen in (26). Average productivity, or absolute advantage, is governed by T_i , which in this model is a CES aggregate of the product-level technology parameters. Since $\theta > \sigma - 1$, average productivity is maximized when the expected number of techniques is equalized across products. This is implicitly assumed to be the case in [Eaton and Kortum \(2002\)](#) but will not be so in general in this model, where the distribution of techniques across products depends on the relative expenditure on different types of research around the world over time.

The second key difference is evident in the presence of the term \tilde{T}_{ii} , which implies that real wages are higher for countries whose production techniques tend to apply to a unique set of products. This is because intermediate products are imperfect substitutes and thus are more valuable in production the less available they are. So, as in [Krugman \(1979\)](#), the factors of production in a country that can produce a unique set of products can extract a share of the surplus generated by their availability in other countries. This expression also demonstrates that, unlike the class of models detailed in [Arkolakis et al. \(2012\)](#), the home trade share is not a sufficient statistic to measure the gains from trade. In autarky, Φ_i^j is proportional to T_i^j , and thus $\tilde{T}_{ii} = 1$. However, as long as techniques are not identically distributed across products for all countries, \tilde{T}_{ii} , will increase as trade barriers are lowered.

2.8 The Value of a Technique

Result 7 above implies that the expected profit for all producers in i from product j is simply a constant share, $1/(1+\theta)$, of expected sales, given by (24). Because the productivity of any producer is independent of when the production technique arrived or to what good it applies, the probability that any particular producer of a given product is the low-cost producer in i is simply the inverse of the expected number of techniques available in i . This implies that, at time t , the expected profit flow to any potential producer of j in i , unconditional on the productivity of the producer, is

$$\Pi_i^j(t) = \frac{1}{1+\theta} \sum_n \frac{(c_i(t)d_{ni})^{-\theta}}{\Phi_n^j(t)} \left(\frac{\Phi_n^j(t)}{\Phi_n(t)} \right)^{\frac{\sigma-1}{\theta}} X_n(t). \quad (29)$$

Thus, the expected discounted present value of future profits (at time t prices) from t onward

based on a technique for producing j in i is

$$V_i^j(t) = P_i^y(t) \int_0^\infty e^{-rs} \frac{\Pi_i^j(t+s)}{P_i^y(t+s)} ds, \quad (30)$$

where r is the interest rate.

3 Equilibrium

I characterize the infinite-horizon world equilibrium in two parts. First, I define a static international trade equilibrium at a point in time, where the technology parameters, $J(t)$ and $\{T_i^j(t)\}$, and the share of the labor force in each country engaged in production are taken as given. Then, I construct a balanced growth path equilibrium in which the world economy is in static equilibrium at every point in time, and the labor shares devoted to the three types of research are constant over time.

3.1 Instantaneous Market Clearing

A static international trade equilibrium is a wage vector $\{w_i\}$ such that, given $J(t)$, $\{T_i^j(t)\}$, and $\{s_i^x(t)\}$, producers of the final good, the composite manufactured good, and individual intermediate goods minimize their per-unit costs subject to (2), (3), and (4), respectively (implying respective costs given by (7), (6), and (8)); aggregate bilateral trade flows are given by (25); and the following conditions are met:

1. Household budget constraint:

$$P_i^y c_i(t) = w_i + \frac{1}{L_i} \left[\sum_n \pi_{ni} \frac{X_n}{1+\theta} - P_i^y Y_i \sum_x s_i^x \right].$$

2. Aggregate resource constraint:

$$Y_i \left(1 - \sum_x s_i^x \right) = L_i c_i, \quad \forall i.$$

3. Composite manufactured good market clearing:

$$Q_i^f + \int_0^{J(t)} Q_i^j dj = Q_i, \quad \forall i.$$

4. Labor market clearing:

$$L_i^f + \int_0^{J(t)} L_i^j dj = L_i, \quad \forall i.$$

5. Balanced trade:

$$X_i = \sum_n \pi_{ni} X_n, \quad \forall i.$$

The term in brackets in the household budget constraint is aggregate profits net of research costs, which uses result 7, above, to express aggregate gross profit as a constant share of aggregate intermediate good sales. The first three conditions and the constant expenditure shares on final and composite manufactured goods that obtain from firm optimization imply that labor income makes up a constant share, $(\gamma + \beta\theta)/((1 - \gamma)(1 + \theta))$, of expenditure on the composite intermediate. Applying this result to the balanced trade condition implies that, in equilibrium,

$$w_i L_i = \sum_n \pi_{ni} w_n L_n, \quad \forall i. \quad (31)$$

Alvarez and Lucas (2007) show that (31) defines a contraction mapping on $\{w_i\}$, and thus there exists a unique static international trade equilibrium.

3.2 Balanced Growth

The values of the set of the technology parameters, $J(t)$ and $\{T_i^j(t)\}$, at every point in time are determined by the history of labor shares devoted to each type of innovative activity in each country. A balanced growth path equilibrium is a set of labor shares, $\{s_i^x\}$, for each country and innovative activity which are constant over time, and for which (a) $J(t)$ is consistent with (14); (b) $\{T_i^j(t)\}$ are given by (19); (c) at every point in time, given $J(t)$ and $\{T_i^j(t)\}$, the set of wages, $\{w_i\}$, is consistent with (31); and (d) at every point in time, the free entry conditions (13) are satisfied. In what follows, the time variable is omitted where ratios are constant over time.

3.3 Technology Along the Balanced Growth Path

Along the balanced growth path, it must be the case that final output in every country grows at a constant rate, g_Y . Equation (14) then implies that $J(t)$ also grows at a constant rate, $g_J = \lambda g_Y$, and that

$$\frac{J}{(Y)^\lambda} = \frac{R^h}{g_J}, \quad (32)$$

where $R^h = \sum_i \alpha_i^h \left(s_i^j \frac{Y_i}{Y} \right)^\lambda$ and $Y(t) = \sum_i Y_i(t)$. Combining this result with (15) implies that

$$\hat{T}_i^{vj}(t) = g_J \frac{R_i^v}{\sum_i R^h} \equiv \kappa_i^v,$$

which further implies that

$$T_i^{vj}(t) = T_{imj}(a^j) = \mathbf{1}_{imj} + \kappa_i^v a^j,$$

where $a^j = t - \underline{t}^j$ is the amount of time since j was invented – i.e., the age of j – and

$$T_w^j(t) = T_w(a^j) = 1 + \bar{\kappa}^v a^j,$$

where $\bar{\kappa}^v = \sum_i \kappa_i^v$. Thus, at every point in time, the stock of original research that applies to a product is simply a function of its age and the country in which it was invented.

Combining all of this with (18) defines the law of motion for T_i^j as a differential equation in the product's age, a , and country of origin, m , given by

$$\dot{T}_{im}(a) = \kappa_i^v + \kappa_i^l(T_w(a) - T_{im}(a)),$$

where $\kappa_i^\ell = gJ \frac{R_i^\ell}{\sum_i (R_i^h + R_i^v)}$, with initial condition given by (16), which has the following solution:

$$T_{im}(a) = T_w(a) - (1 - \mathbf{1}_{im})e^{-\kappa_i^\ell a} - \frac{\bar{\kappa} - \kappa_i^v}{\kappa_i^\ell}(1 - e^{-\kappa_i^\ell a}). \quad (33)$$

Because $J(t)$ grows at a constant rate, the share of products that are of age a at any point in time is constant and given by an exponential distribution with parameter gJ , and the stock of products that were invented in each country is also a constant share of $J(t)$, given by $\eta_i = \frac{R_i^h}{\sum_i R_i^h}$. Thus, the probability that any given product is of age a and was invented in m is given by

$$f(a, m) = \eta_m gJ e^{-gJ a}. \quad (34)$$

Equations (33) and (34) fully characterize the distribution of research that has been applied across all countries and products at a point along the balanced growth path. The world stock of original research which applies to a particular product simply increases linearly over time with a slope equal to the total flow of world vertical research relative to the measure of products. The amount of research, both original and learned, that has been applied in country i lies below this frontier. The second term in (33) captures the initial advantage that the inventor of a product receives, which erodes over time as researchers in other countries learn about the production techniques developed by the inventor. The third term measures the the amount of original vertical research done outside a country that has not yet been learned by researchers there. This term grows over time as the world frontier expands, but it grows more slowly than the frontier – reflecting the distance-to-the-frontier property of learning – meaning that the *share* of worldwide original research that has been applied in i grows through learning as a product ages. The production of original vertical research increases both the stocks of domestic and worldwide original research, so this effect is subsumed in $T_w(a)$.

3.4 Equilibrium Growth Rates

Applying the results from the previous section to (31) implies that relative wages across countries are constant over time, so, choosing w_1 as the numeraire, from (22) it is straightforward to show

that the price of the of the composite intermediate good in n can be expressed as

$$P_n(t) = J(t)^{\frac{-1}{\beta(\sigma-1)}} P_n \quad (35)$$

where

- $P_n = \hat{\gamma} \Phi_n^{-\frac{1}{\theta}}$,
- $\Phi_n = \left(\sum_m \int_0^\infty \Phi_{nm}(a)^{\frac{\sigma-1}{\theta}} f(a, m) da \right)^{\frac{\theta}{\sigma-1}}$,
- $\Phi_{nm}(a) = \sum_i (c_i d_{ni})^{-\theta} T_{im}(a)$, and
- $c_i = \tilde{\beta} w_i^\beta P_i^{1-\beta}$.

Substituting (35) into (7) makes clear that the price of final goods in every country falls at a constant rate given by

$$g_c \equiv -\frac{\dot{P}_i^y}{P_i^y} = \frac{1-\gamma}{\beta(\sigma-1)} g_J, \quad \forall i, \quad (36)$$

which is also the growth rate of per-capita final output and consumption since nominal income per worker (relative to w_1) is constant along the balanced growth path. Since $g_Y = g_c + g_L$, this implies that

$$g_J = \frac{\lambda}{1 - \lambda \frac{1-\gamma}{\beta(\sigma-1)}} g_L,$$

given that $\lambda \frac{1-\gamma}{\beta(\sigma-1)} < 1$.¹⁴

3.5 Equilibrium Research

From, (25), (29), and (35), it is straightforward to show that, along the balanced growth path, the ratio of expected profit derived from a given technique to total profit *per product* is stationary and given by

$$\frac{\Pi_{im}(a)}{\Pi_i/J} = \frac{1}{T_i} \sum_n \frac{1}{\tilde{T}_{ni}} \frac{X_{ni}}{X_i} \left(\frac{\Phi_{nm}(a)}{\Phi_n} \right)^{\frac{\sigma-1}{\theta}-1}, \quad (37)$$

where T_i and \tilde{T}_{ni} are the stationary values of their counterparts defined by (26) and (27), respectively, given by

$$T_i = \left(\sum_m \int_0^\infty T_{im}(a)^{\frac{\sigma-1}{\theta}} f(a, m) da \right)^{\frac{\theta}{\sigma-1}}$$

¹⁴Otherwise, as in other models of “semi-endogenous” growth, such as Jones (1995), Kortum (1997), and Segerstrom (1998), with population growth, the growth rate increases over time. I have chosen a model of this form purely for simplicity as this paper is primarily concerned with the cross-country distribution of income along the balanced growth path, not its growth rate. In a related framework in which the R&D production function admits fully endogenous growth, Howitt (2000) shows that, as in this paper, countries grow at the same constant rate in the long run, and differences in R&D output result in differences in relative income per worker.

and

$$\tilde{T}_{ni} = \sum_m \int_0^\infty \frac{T_{im}(a)}{T_i} \left(\frac{\Phi_{nm}(a)}{\Phi_n} \right)^{\frac{\sigma-1}{\theta}-1} f(a, m) da.$$

This result, along with (32), implies that, along the balanced growth path, (30) can be rewritten as

$$V_{im}(a, t) = \frac{1-\gamma}{1+\beta\theta} \frac{gJ}{R^h} \left(\frac{Y_i}{Y} \right)^\lambda P_i^y(t) Y_i(t)^{1-\lambda} V_{im}(a), \quad (38)$$

where

$$V_{im}(a) = \int_t^\infty e^{-\nu s} \frac{\Pi_{im}(a+s)}{\Pi_i/J} ds,$$

and $\nu = r - (1-\lambda)g_Y$.¹⁵ To understand equation (38), note that it simply decomposes the expected value of a given production technique at a point along the balanced growth path into the product of the ratio of total profit to GDP, GDP per product, and the discounted expected value of future profit flows from the technique relative to total profit per product, where the appropriate discount rate takes into account that sales per product (in period t prices) grows over time.

The form of (38) implies that the expected value of obtaining a production technique through one of the three types of research activity can be expressed as

$$V_i^x(t) = \frac{1-\gamma}{1+\beta\theta} \frac{gJ}{R^h} \left(\frac{Y_i}{Y} \right)^\lambda P_i^y(t) Y_i(t)^{1-\lambda} V_i^x.$$

The stationary components of these expressions are given by the following:

$$\begin{aligned} V_i^h &= V_{ii}(0), \\ V_i^v &= \sum_m \int_0^\infty V_{im}(a) f(a, m) da, \\ V_i^\ell &= \sum_m \int_0^\infty V_{im}(a) \frac{T_w(a) - T_{im}(a)}{T_w/J} f(a, m) da. \end{aligned} \quad (39)$$

The first of these expression is the most straightforward. The value of inventing a new product is equal to the flow of profit of a product of age 0 that was invented in the country of the producer. Because vertical innovation is undirected, the expected value of a new technique resulting from vertical innovation is simply the expected value of $V_{im}(a)$ evaluated over the joint probability distribution of a and m . The value of a technique acquired through learning is similar, except that, from (17), the probability of obtaining a useable technique that applies to a particular product is proportional to the world stock of original research pertaining to the product that has not yet been utilized in i .

These results imply that, along a balanced growth path, the free entry conditions (13) are given

¹⁵The static equilibrium conditions imply that $\Pi_i/P_i^y Y_i = (1-\gamma)/(1+\beta\theta)$.

by

$$\frac{1 - \gamma}{1 + \beta\theta} \frac{g_Y}{R^h} \alpha_i^x (s_i^x)^{\lambda-1} \left(\frac{Y_i}{Y} \right)^\lambda V_i^x = 1, \quad \forall i, x. \quad (40)$$

These conditions define the equilibrium values of $\{s_i^x\}$ along the balanced growth path.¹⁶

3.6 Innovation Incentives

Before solving for the balanced growth path equilibrium numerically, it is useful to gain some intuition into the forces that drive countries' innovation decisions and shape the world income distribution. A country in this model is defined by its population size, $L_i(t)$, its level of productivity in research of all types, α_i , and its barriers to trade with the other countries of the world, d_{ni} . Considering a few special cases will help make clear how each affect the decisions of potential entrants.

3.6.1 A Country in Isolation

First, consider the case in which a country is completely isolated from the rest of world, and note that this is equivalent to a case in which the value of either α_i or L_i , relative to other countries' values, becomes arbitrarily large. In this case, trade flows (relative to sales in i) are equal to zero, and $T_w(a) = T_{ii}(a)$, so learning is not profitable for entrants in i . Equation (37) reduces to

$$\frac{\Pi_{im}(a)}{\Pi_i/J} = \frac{T_{im}(a)^{\frac{\sigma-1}{\theta}-1}}{T_i^{\frac{\sigma-1}{\theta}}}.$$

With probability 1, a given product was invented in i , so only the case in which $m = i$ is relevant here, and $\kappa_i^v = \alpha^v (s_i^v)^\lambda / \alpha^h (s_i^h)^\lambda$ – which implies that both $T_{ii}(a)$ and T_i , and thus V_i^v , are independent of both α_i and L_i . Also, in this case, $Y_i = Y$, and R^h is proportional to α_i . Together, through (40), this implies that s_i^h and s_i^v are independent of α_i and L_i , and $s_i^\ell = 0$. Intuitively, a greater value of α_i or L_i increases the return to research by, respectively, increasing the likelihood of success of research or the size of the market for a new product or technique. However, it also implies a greater number of products and techniques with which a potential entrant would compete for revenue. When a country is large or isolated, the number of products increases by exactly as much as the probability of research success or aggregate expenditure in intermediate goods, keeping the number of techniques per product constant. Thus, the two forces exactly offset, and α_i and L_i have no effect on equilibrium research shares. This result does not depend on the values of d_{ni} or α^ℓ , reflecting the fact that international flows of goods and ideas are relatively unimportant for countries that can produce relatively large amounts of research.

¹⁶While this does not constitute a proof of the existence and uniqueness of a balanced growth path equilibrium, in practice, for many sets of parameter values, it was always possible to find an equilibrium computationally, and a wide variety of starting values of the endogenous variables always led to the same equilibrium. In fact, (40) appears to implicitly define a contraction mapping on $\{s_i^x\}$.

3.6.2 Access to Foreign Product Knowledge

In general, however, access to foreign products and ideas is important in determining innovation expenditure in a given country. To see how access to the knowledge of the existence of products that were invented abroad matters, consider a country for which α_i or L_i is infinitely small relative to other countries in a world in which $d_{ni} = \infty$ for all $n \neq i$ and $\alpha^\ell = 0$ so that, as above, there are no trade flows and no learning about vertical research. In this case, a measure zero share of products were invented in i , κ_i^v is proportional to $\alpha_i(Y_i/Y)^\lambda$, and $\Pi_{im}(a)/(\Pi_i/J)$, where $m \neq i$, and thus V_i^v is inversely proportional to the same term for all a . On the other hand, $\Pi_{ii}(0)/\Pi_i$, which determines the value of horizontal research, is inversely proportional to $(\alpha_i(Y_i/Y)^\lambda)^{(\sigma-1)/\theta}$. Thus, a marginal increase in α_i or L_i increases the value of horizontal relative to vertical research.

Intuitively, a small country does not affect the measure of products available in equilibrium, so a proportional increase in α_i or L_i increases the number of techniques per product in i . The increased likelihood of research success or market size is exactly offset by the fall in the probability of being the low-cost producer for a technique to produce a randomly drawn product. For a newly invented product, however, the producer is guaranteed to be the low-cost producer, and because products are imperfect substitutes, the increase in probability of research success or market size is only partially offset by the shift in expenditure toward other products, for which the expected price is lower.

3.6.3 Access to Trade

To see how access to foreign goods markets affects research incentives, suppose that in the previous case, $d_{ni} = 1$, so that trade is frictionless. In this case, $\Phi_{nm}(a) = \Phi_m(a)$, for all n , and

$$\frac{\Pi_{im}(a)}{\Pi_i/J} = \frac{\Phi_m(a)^{\frac{\sigma-1}{\theta}-1}}{\hat{T}_i},$$

where $\hat{T}_i = \sum_m \int_0^\infty T_{im}(a) \Phi_m(a)^{\frac{\sigma-1}{\theta}-1} f(a, m) da$. Since $\Phi_m(a)$, for $m \neq i$, is negligibly affected by α_i and L_i , this expression is inversely proportional to $\alpha_i(Y_i/Y)^\lambda$ for all a , and thus so is V_i^v . However, for a newly invented product,

$$\frac{\Pi_{ii}(0)}{\Pi_i} = \frac{c_i^{-(\sigma-1)}}{\hat{T}_i c_i^{-\theta}}.$$

Since under free trade, c_i is proportional to w_i , horizontal innovation becomes relatively more profitable for higher wage countries. From (28), the wage in country i under free trade can be expressed as

$$\frac{w_i}{P} = \left(\frac{\hat{T}_i}{L_i} \right)^{\frac{1}{1+\beta\theta}},$$

which implies that w_i is increasing in $\alpha_i(Y_i/Y)^\lambda/L_i$. This means that the relative returns to horizontal research are increasing in α_i and decreasing in L_i .

As in the previous case, a proportional increase in α_i or L_i increases the number of techniques per product in i . With trade, it is worldwide expenditure on a product that matters for expected profits, and because i is small relative to the rest of the world, a proportional increase the number of techniques in i has no effect on the expected level of expenditure on products for which the rest of the world has accumulated techniques, and only leads to a relative increase in the the probability that the low-cost producer is from i . Thus, expected expenditure per product, and thus the returns to vertical research, are unaffected by marginal changes in α_i or L_i . For a product that has just been invented in i , however, the inventor is guaranteed to be the low-cost producer and is “large” relative to the market for that product. As a result, expected profits depend on local production costs, which determine relative expenditure on the product through the demand elasticity, $\sigma - 1$. An increase in α_i increases the probability of successfully inventing a product, but also leads to an increase in demand for labor in i and, in turn, wages. Because $\theta > \sigma - 1$, the net effect on expected profits is positive. To understand why, first note that, in aggregate, the net effect on profits per technique of an increase in α_i , including the effect on wages, is zero. Second, for all but a measure zero set of products, which were invented in i , wages only affect the probability that the low-cost producer is in i , with an elasticity given by θ . Thus, the effect of a change in α_i on expected profits, in the absence of the effect on wages, is equal the effect of wages, with an elasticity of θ . So, the net effect for the inventor of a product depends on the change in wages with an elasticity of $\theta - (\sigma - 1)$.

3.6.4 Access to Foreign Vertical Research

To understand how the returns to learning depend on country characteristics, note that the value of a technique obtained through learning is very similar to that from vertical research, with two key differences. First, because the world technological frontier grows over time, techniques from learning will tend to apply to older products, and second, the likelihood that a technique is useable is increasing in a country’s average distance to the technological frontier. The similarity means that α_i and L_i have a similar effect on the returns to learning as on the returns to vertical research, and both differences amplify this effect. The fact that techniques obtained through learning apply to older products means that an increase in techniques per product in a small country due to an increase in α_i or L_i will decrease the return to learning to a greater extent than it will vertical research. And, a country’s distance to the technological frontier for every product is mechanically decreasing in α_i and L_i . Thus, countries with larger relative values of α_i and L_i devote a smaller share of output to learning.

The existence of learning also affects the returns to the other innovative activities. Learning, like vertical research, increases the number of techniques available per product in each country, and because the arrival rate of new techniques is proportional to the world technological frontier, it further skews the distribution of techniques toward older products. This implies that an increase in s_i^ℓ decreases the expected return to all three innovative activities, but that it affects the return to vertical research, which applies to older products, by more than it does the return to horizontal research. Further, with finite trade barriers, increased learning in one country affects research

incentives in all countries.

This is clear from (37), where the profit flow from a given technique at a point in time is decreasing in $\Phi_{nm}(a)$ with an elasticity of $(\sigma - 1)/\theta - 1$. The existence of learning implies that $\Phi_{nm}(a)$ grows faster over time by increasing the number of techniques available in every country but also by increasing the relative number of techniques in lower-wage countries. This decreases expected profit flows from techniques applying to older products. The magnitude of this effect is smaller the greater is σ and smaller is θ because a larger value of $(\sigma - 1)/\theta$ implies that relatively more expenditure is devoted to products for which $\Phi_{nm}(a)$ is relatively large, since a larger value of σ implies that products are relatively more substitutable, and a smaller value of θ implies that an additional technique is more likely to lead to a lower cost of producing a product.

3.6.5 Trade Barriers

Finally, to see the effect of the level of trade barriers, and not simply the discrete jump from autarky to free trade, on research incentives, consider how a proportional global decrease in all trade costs would affect expected profit flows derived from a technique. To do this, write $d_{ni} = \bar{d}_{ni}d$ for all $n \neq i$, and consider the partial elasticity of $\Pi_{im}(a)$ with respect to d (holding X_n constant), which is given by

$$\frac{\partial \ln \Pi_{im}(a)}{\partial \ln d} = \theta \hat{X}_{iim}(a) - \sum_n \hat{X}_{nim}(a) [\theta \pi_{nmm}(a) + (\sigma - 1)(\pi_{nn} - \pi_{nmm}(a))],$$

where $\hat{X}_{nim}(a) = E[X_{nim}(a)] / \sum_n E[X_{nim}(a)]$ is the share of expected sales (and profits) for a given product that come from sales in n . This expression illustrates the opposing forces of trade costs on firms' profits: higher costs of reaching foreign markets and a decrease in competition in the domestic market from foreign firms. The first term represents the increase in the likelihood that a firm becomes the low-cost provider to the home market due to the increased costs of importing. The first term in brackets represents the fall in this probability in foreign markets as it becomes relatively cheaper for consumers there to buy from local firms, and the second term in brackets represents the shift in expenditure in every market toward products for which local firms have a comparative advantage.

To see how research incentives are influenced, consider this elasticity for the special case of a product newly invented in i , given by

$$\frac{\partial \ln \Pi_{ii}(0)}{\partial \ln d} = (\sigma - 1) \hat{X}_{iii}(0) - (\sigma - 1) \sum_n \hat{X}_{nii} \pi_{nn}.$$

Notice that the terms multiplied by θ have cancelled out. This is because a producer of a newly invented product has no direct competitors worldwide – i.e. $\pi_{nii}(0) = 1$. This means that trade barriers only affect a firm's expected profits through the reallocation of expenditure in each market toward the products in which local firms have a comparative advantage. Notice that, because $\hat{X}_{nii}(0) \geq \hat{X}_{nii}(a)$, $\hat{X}_{iii}(0) \leq \hat{X}_{iii}(a)$, and all terms that disappear from the general elasticity

formula are negative, in general $\partial \Pi_{ii}(0)/\partial d \leq \partial \Pi_{im}(a)/\partial d$. Thus, a rise in trade barriers makes horizontal research relatively less profitable.

3.6.6 Summary

All together, this implies that a greater value of α_i or L_i relative to the rest of the world implies that a country will devote relatively less of its output to learning and more to horizontal research. In the absence of learning, it also implies a smaller share devoted to vertical research. However, the returns to learning fall more rapidly than the returns to vertical research as these parameters increase, so it is possible for this relationship to be reversed when researchers can choose among all three research activities. Access to international trade increases the relative return to horizontal research in higher wage countries, which implies that relative expenditure on horizontal research is increasing in α_i and decreasing in L_i as trade barriers fall. And, in the absence of prohibitive trade barriers, greater worldwide expenditure on learning implies a lower return to vertical relative to horizontal research, meaning that a greater level of worldwide learning leads to less expenditure on vertical research both in absolute terms and relative to horizontal research.

3.7 The Equilibrium Cross-Country Income Distribution

To see how income per worker in a given country is affected by these forces, note that, based on (28), along the balanced growth path, income per worker can be expressed as

$$\frac{Y_i(t)}{L_i(t)} = \left(\frac{T_i}{\pi_{ii}} \right)^{\frac{1-\gamma}{\beta\theta}} \tilde{T}_{ii}^{\frac{1-\gamma}{\beta\theta}} J(t)^{\frac{1-\gamma}{\beta(\sigma-1)}} \quad (41)$$

where

$$J(t) = \left(\frac{R^h}{gJ} \right)^{\frac{g_Y}{g_L}} \left[\sum_i \left(\frac{T_i}{\pi_{ii}} \right)^{\frac{1-\gamma}{\beta\theta}} \tilde{T}_{ii}^{\frac{1-\gamma}{\beta\theta}} L_i(t) \right]^{\frac{g_J}{g_L}}.$$

GDP per worker – which, in this model is equivalent to measured TFP – can be broken down into three terms.¹⁷ The first two terms are the stationary counterparts of the those in (28), and the third term makes up the dynamic component of GDP growth.

As discussed in Section 2.7.2, T_i determines average productivity across all products and is increasing in the average number of techniques per product in i and the degree to which these techniques are evenly distributed across products. A lower home trade share increases *measured* productivity by allowing a country to specialize in the products for which it is relatively more productive. And, a higher value of \tilde{T}_{ii} reflects the additional surplus that producers in i obtain from the ability to provide products for which there are few competing techniques in other countries.

The third term summarizes the benefits to all countries in the world from access to a larger measure of intermediate inputs, holding constant the distribution of techniques across products. The expression for $J(t)$ shows that this effect can be further broken down into the effect of the

¹⁷Constant terms are suppressed in this expression.

Table 1: Baseline Parameter Values

Symbol	Value	Target
r	0.07	Rate of return on equity
g_L	0.015	World population growth rate
γ	0.33	Ratio of labor compensation to gross output in non-tradable sectors
β	0.24	Ratio of labor compensation to gross output in manufacturing
σ	4.3	GDP per capita growth rate
θ	4.1	Trade cost elasticity
λ	0.69	Average R&D expenditure
α^ℓ	2.69	Average R&D expenditure among G5 countries
α^h	0.88	Relative number of firms engaged in product innovation

worldwide share of resources that are devoted to horizontal research, which determines the measure of products per worker, and a term related to the stationary component of total world output. This final term arises due to the fact that final output is used in research, which implies that any change that increases world output for a given measure of products is amplified by indirectly increasing gross research output, given the shares of output devoted to research.

4 Quantitative Analysis

This section presents the properties of the general equilibrium of the model along its balanced growth path. The first part of this section discusses the choice of exogenous parameter values. The second presents the model’s quantitative predictions for cross-country income levels, research expenditure, and product-level trade flows and compares these predictions to available data. The third section presents the results of several counterfactual experiments designed to measure the contribution of the access to foreign goods and technologies to research incentives and the world distribution of income.

4.1 Parameter Values

Table (1) lists the values of all the parameters that are common across countries. I set the real interest rate, r , to 7%, approximately the value of the long-run average real return on equity found in Mehra and Prescott (1985) and somewhat above typical estimates of the risk-free rate, as it determines the annual return to R&D expenditure. I set the population growth rate to 1.5%, which is approximately the average annual growth rate of the world population from 1985 through 2005 as reported in the Penn World Tables. The labor endowments are taken from the Penn World Tables. The remaining parameters, described below, are chosen so that the model is consistent with salient features of the data. The data are compiled from multiple sources. All data are from 2006 where available. Details are in Appendix B.

4.1.1 Factor Shares

The parameters of the Cobb-Douglas production functions for final and intermediate goods, γ and β , respectively, determine the share of expenditure on labor and the composite manufactured good by producers in each sector. Therefore, I use data on the the share of labor compensation relative to total output in manufacturing and the non-tradable sectors to set the values of γ and β , respectively.¹⁸

Total labor compensation in manufacturing as a share of gross manufacturing output, taken from the OECD’s STAN database, ranges from about 0.1 to 0.25, with values generally between 0.15 and 0.25 for the western European and North American countries and the value for the United States equal to 0.19. [Gollin \(2002\)](#) argues that labor compensation is often underestimated in less developed countries, so the U.S. value seems like a reasonable estimate for the labor share of manufacturing output. In the model, this share is equal to $\beta\theta/(1 + \theta)$, so I set $\beta = 0.19(1 + \theta)/\theta$.

The values of labor compensation as a share of total output in the non-tradeable sectors ranges between approximately 0.2 and 0.35, and generally between 0.25 and 0.35 for the Western European and North American countries, with the U.S. on the high end of the range at 0.34.¹⁹ As a result, I use a value of 0.33 for γ , slightly below the U.S. value but still on the high end of advanced economies.

4.1.2 Elasticities

The parameter θ governs the degree of dispersion of productivity across production techniques and also serves as the (partial) elasticity of bilateral trade flows with respect to trade costs. Many attempts have been made to estimate this elasticity. [Alvarez and Lucas \(2007\)](#) surveys many of these studies and argues for a value in the range $[4, 10]$, with their preferred value being $\theta = 6.7$. [Eaton and Kortum \(2002\)](#) estimate values in the range $[3.6, 12.8]$, with their preferred estimate being $\theta = 8.3$. However, [Waugh and Simonovska \(2013\)](#), argue that their estimators are biased, and using international price data and an unbiased method based on the model of [Eaton and Kortum \(2002\)](#), estimate values in the range $[2.8, 4.5]$, with a benchmark estimate of $\theta = 4.1$. Because the framework of [Waugh and Simonovska \(2013\)](#) is much more closely related to this one than the studies on which [Alvarez and Lucas \(2007\)](#) base their choice of θ , I set $\theta = 4.1$ as a benchmark. However, the quantitative results, especially the effect of α^ℓ on the world income distribution, turn out to be sensitive to the value of θ , so I also consider $\theta = 6.7$ as an alternative value.²⁰

¹⁸[Alvarez and Lucas \(2007\)](#) consider labor in their model to be “labor plus capital” and use data on total value added to fix the analogues of these parameters in a static model. However, in a growth framework, the capital-labor ratio is endogenous and variable over time, so I define labor in the more traditional way. To avoid the added complexity of introducing a third factor of production, I thus interpret the composite manufactured good as “intermediates-plus-capital”, which implicitly assumes that all capital depreciates after one period.

¹⁹The set of non-tradeable sectors corresponds to all industries that fall under the 2-digit ISIC (Revision 3) codes 40-99, which includes electricity, gas, and water supply; construction; and all services.

²⁰Most of these estimations use data that is aggregated at least to the industry level. However, in this model, such a practice is problematic as the the trade cost elasticity of aggregated trade flows lies between $\sigma - 1$ and θ . The

The elasticity of substitution between individual intermediate goods in the production of the composite manufactured good, σ , governs the magnitude by which an increase in $J(t)$ leads to an increase in output of the composite manufactured good and, in turn, final output. As a result, it plays a role in determining the steady state growth rate of per capita income, g_c , given by (36). Therefore, I choose σ so that, given values for β , γ , and λ , the growth rate of per capita income in steady state is equal to 2.2%. Because the value of β depends on θ , so does the value of σ . When $\theta = 4.1$, this implies a value of $\sigma = 4.5$, while $\theta = 6.7$ implies $\sigma = 5.1$. These values are well within the range estimated by Broda and Weinstein (2006) based on US import data.

4.1.3 R&D Productivity Parameters

The country-specific research productivity parameters, $\{\alpha_i\}$, are chosen to match gross manufacturing output for each country. Because only relative productivity across both countries and research activities matters in the general equilibrium, one of these productivity parameters can be normalized, so I set $\alpha^v = 1$. The remaining parameters – α^h , α^ℓ , and λ – are chosen jointly so that the predictions of the model match three moments from data on research and development expenditure.

The first two moments are based on data on total R&D expenditure as a percentage of GDP. The first is the average level of R&D expenditure across all countries for which data is available, and the second is the average over the G5 countries, respectively 1.4 and 2.5%. In the model, I consider R&D expenditure to include expenditure on horizontal and vertical research but not learning, which would plausibly show up in the data as engineering or other non-production but also non-R&D expenses. The third moment, based on data from the Community Innovation Survey available from Eurostat, is the average ratio of the number of enterprises engaged in product innovation to the number engaged in process innovation, which is equal to 0.94 over all countries for which data is available. In the model, this ratio is equivalent to the ratio of total horizontal to vertical research expenditure for a given country.

While each of these coefficients influences all three moments, each primarily affects a single moment in an intuitive way. Given the normalization of α^v , α^h controls the relative return to horizontal versus vertical research, with a higher value of α^h implying a higher equilibrium value of s_i^h/s_i^v . Similarly, α^ℓ governs the relative return to learning, so a higher value tends to lower measured R&D expenses for all countries. A higher value of λ disproportionately increases the return to research by countries that are already heavily engaged in research, which tend to be large, productive countries, while at the same time increasing the world technological frontier and disproportionately increasing the return to learning in countries that are far from it. Thus, a higher value of λ tends to increase R&D expenditure in the G5 countries relative to the rest of the world.

estimation of Waugh and Simonovska (2013) is based on the relationship $\ln\left(\frac{X_{ni}/X_n}{X_{ii}/X_i}\right) / \ln\left(d_{ni} \frac{P_i}{P_n}\right) = -\theta$, while in this model $\ln\left(\frac{X_{ni}/X_n}{X_{ii}/X_i}\right) / \ln\left(d_{ni} \frac{P_i}{P_n}\right) = \ln\left(\frac{\bar{T}_{ni}}{\bar{T}_{ii}}\right) - \theta$. However, using the numerical results of the model for a given value of θ in the range discussed above, it turns out that the bias imposed by ignoring the additional term is quite small.

4.1.4 Trade Costs

The final set of parameters to be fixed are the bilateral iceberg trade costs, d_{ni} . I parameterize trade costs as a function of observable variables suggested by the gravity literature. Specifically, I assume that

$$\ln d_{ni} = d_k + bord_{ni} + lang_{ni} + col_{ni} + rta_{ni} + ex_i, \quad (42)$$

where d_k is the effect of the distance between n and i lying in one of six distance intervals; $bord_{ni}$ is the effect of countries n and i sharing a common border; $lang_{ni}$ the effect of sharing a common language; col_{ni} the effect of having a colonial relationship; rta_{ni} the effect of n and i being part of a regional trade agreement; and ex_i an exporter-specific border cost.²¹

To estimate the parameters of d_{ni} , I take advantage of the gravity-like structure of the model-implied relationship between bilateral trade flows and trade costs. However, unlike [Eaton and Kortum \(2002\)](#) and related papers, this model does not imply a standard gravity relationship between aggregate bilateral trade flows and trade costs due to the presence of the \tilde{T}_{ni} term in (25), which makes it impossible to disentangle the effect of trade costs from that of the interaction among countries' patterns of comparative advantage. Therefore, their estimation methods are not appropriate here.

The relationship in this model between product-level trade flows and trade costs, given by (21), on the other hand, does not suffer from this issue. However, using this as the basis of a trade cost estimation does entail some unique issues which must be addressed. The first is that there is no data on trade flows at the level of disaggregation implied by the model, i.e. the level at which products are so narrowly defined that no country would ever purchase the same product from more than one source. The most disaggregated bilateral trade data available for a large number of countries is that from the U.N.'s Comtrade database classified at the 6-digit level of the Harmonized System (HS-6). In order for estimates based on this data not to be subject to the same criticism as those based on aggregate data, it must be the case every country's value of T_i^j must not vary across the products that lie within each HS-6 category. While this is a strong assumption, it is undoubtedly much more plausible for a set of products within an HS-6 category than for all products that make up the aggregate data. Thus, I proceed in this direction.

Let us assume that, at a point in time, the product space is divided into discrete categories, denoted $b = 1, \dots, B$, which correspond to HS-6 categories, in which it is approximately the case that $T_i^j = T_i^b$ for all $j \in b$. Then, denoting by M_n^b imports to n in category b , (21) implies that the value of exports from i to n in category b are equal to

$$X_{ni}^b = \frac{T_i^b (c_i d_{ni})^{-\theta}}{\tilde{\Phi}_n^b} M_n^b, \quad (43)$$

where $\tilde{\Phi}_n^b = \sum_{i \neq n} T_i^b (c_i d_{ni})^{-\theta}$.

²¹I assume that this effect is exporter-specific, rather than importer-specific, following [Vaugh \(2010\)](#), which argues that this specification is more consistent with data on the prices of tradable goods.

Table 2: Trade Cost Coefficient Estimates

Variable	Coefficient	Std. Err.
<625 km	-5.03	-
625 - 1,250 km	-5.53	0.13
1,250 - 2,500 km	-5.96	0.18
2,500 - 5,000 km	-6.34	0.27
5,000 - 10,000 km	-7.24	0.29
>10,000 km	-7.52	0.34
Shared Border	0.36	0.10
Common Language	0.24	0.06
Colonial Ties	0.11	0.08
Trade Agreement	0.66	0.11

Notes: Number of observations: 37,663,020. Standard errors are clustered by exporter. Reported coefficient estimates are equal to $-\theta b$, where \hat{b} corresponds to the parameters of (42). Distance coefficients are equal to $-\theta(d_k + \frac{1}{N} \sum_i ex_i)$.

This brings us to the second issue. The most conventional approach to estimating (43) would be to proxy for the terms $T_i^b c_i^{-\theta}$ and $M_n^b / \tilde{\Phi}_n^b$ using exporter-product and importer-product fixed effects. However, the dataset contains 81 countries and 4,702 manufacturing product categories, which implies that the estimation would involve 761,724 dummy variables, which is computationally infeasible. Instead, I use the method of French (2014), which uses data on total imports and exports by country and product category to eliminate the need for fixed effects.²²

To apply this method here, define $E_i^b = \sum_{n \neq i} X_{ni}^b$ to be total exports of i in category b . It is straightforward to show that X_{ni}^b can be expressed as a function of only trade costs and the set of total exports and imports in category b of all countries by the following system of equations:

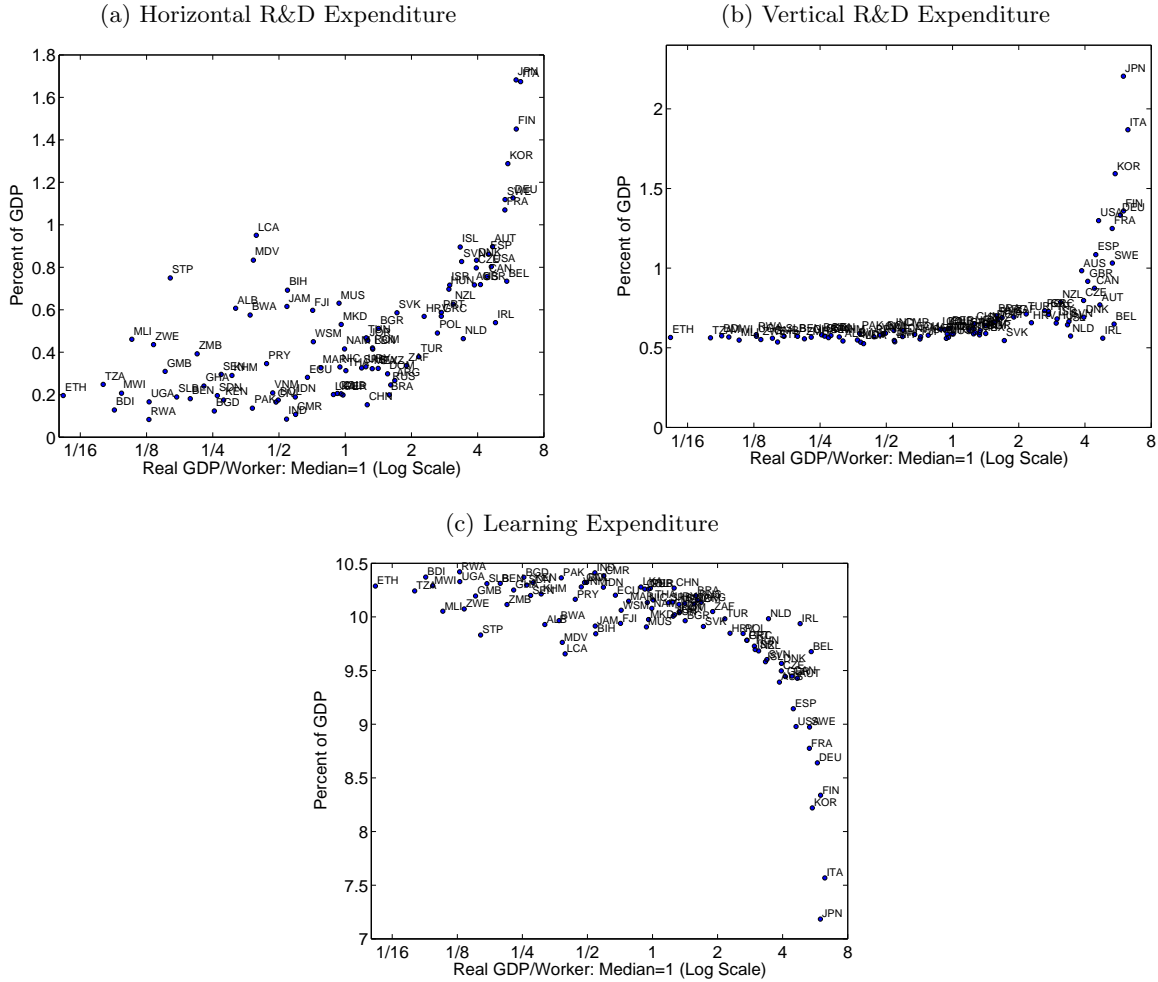
$$\begin{aligned}
X_{ni}^b &= \frac{E_i^b M_n^b}{E^b} \frac{d_{ni}^{-\theta}}{\tilde{\Phi}_n^b \Psi_i^b} + \varepsilon_{ni}^b, \\
\Psi_i^b &= \sum_{n \neq i} \frac{d_{ni}^{-\theta}}{\tilde{\Phi}_n^b} \frac{M_n^b}{E^b}, \\
\tilde{\Phi}_n^b &= \sum_{i \neq n} \frac{d_{ni}^{-\theta}}{\Psi_i^b} \frac{E_i^b}{E^b},
\end{aligned} \tag{44}$$

where $E^b = \sum_i E_i^b$, and ε_{ni}^b is appended to reflect measurement error.

The final issue is that, using only data on product-level imports and exports and not domestic trade, the parameter ex_i is not identified. This is because only relative trade costs matter in determining trade flows in (44), and thus scaling the cost of exporting up or down for a given exporter will not affect its bilateral trade flows *conditional on total exports*. Since data on domestic trade flows are not available at the product level, as in French (2014), I rely on data on aggregate

²²This method is related to the aggregate nonlinear least squares estimator of Anderson and van Wincoop (2003).

Figure 1: R&D Expenditure and Income per Worker




domestic trade flows, taken as total manufacturing output minus exports, to identify ex_i , setting ex_i equal to the value necessary for the value of X_{ii} predicted by the model to match the data, given all the other parameter values.

Following Santos Silva and Tenreyro (2006), I estimate (44) via Poisson pseudo maximum likelihood. Coefficient estimates are presented in Table 2. The results are generally consistent with the gravity literature. Trade flows are found to be decreasing in distance and increasing in shared borders, language, colonial relationships, and trade agreements.

4.2 Numerical Results

Given parameter values, it is possible to compute equilibrium research shares, trade flows, and the world distribution of income. Details of the computation procedure are in Appendix C. Figure 1 depicts the equilibrium shares of income devoted to each innovative activity. Consistent with the discussion in Section 3.6, high-income countries devote the largest shares of output to horizontal

Figure 2: Real GDP per Worker: Model vs. Data



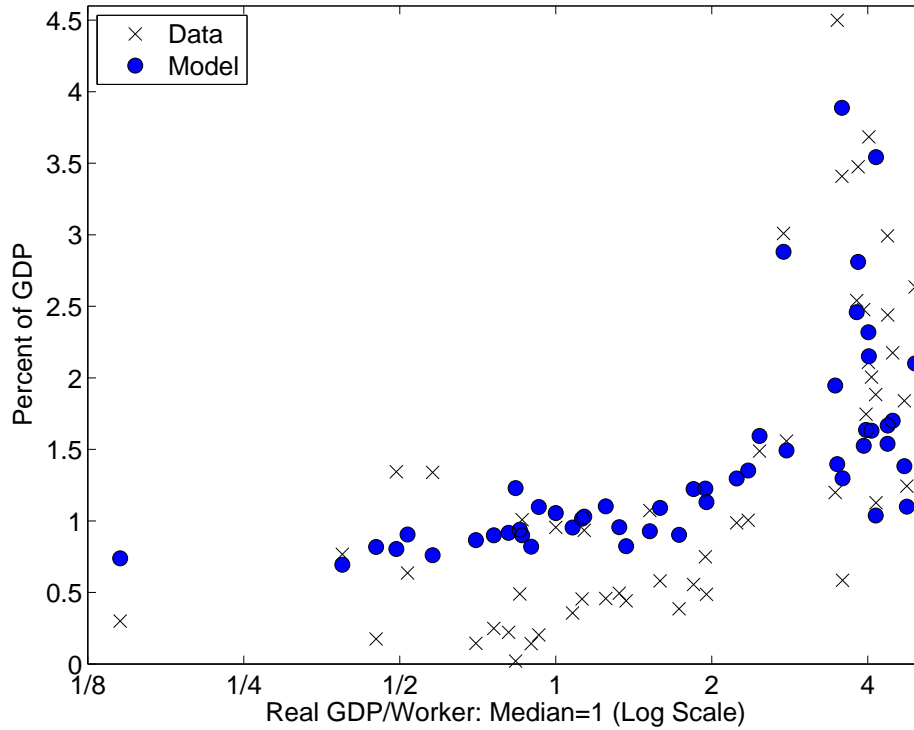
research, while low-income countries spend a much larger share on learning. It is also evident that, given the parameter values, returns to learning increase much more rapidly than returns to vertical research as countries move farther from the technological frontier, so that high-income countries also devote a much larger share of income to vertical research.

4.2.1 Income and R&D

Figure 2 plots the levels of real GDP per worker predicted by the model against those in the data, along with the 45 degree line. It is evident that the model does quite well in matching real income levels despite the fact that neither nominal GDP nor price levels were targeted when choosing parameters. The variance in log income per worker is 1.54 in the model versus 1.55 in the data, and the 90/10 ratio of income per worker is 30.4 in the model versus 25.3 in the data.

Figure 3 illustrates the degree to which the model captures the relationship between income and R&D expenditure present in the data. It plots both the actual model-predicted values of the share of GDP devoted to R&D against real GDP per worker for 2006. As above, I consider expenditure on horizontal and vertical research, but not learning, in the model to be R&D expenditures. While the model parameters were chosen to match average R&D expenditure across all countries and among the G5 countries, the degree to which the model's predictions match the shape of the relationship between income and R&D expenditure can be viewed as a test of the model's validity. As can be seen, the shape of the relationship predicted by the model is quite similar to that in the data,

Figure 3: R&D Expenditure and GDP: Model and Data



with R&D expenditure remaining relatively flat with rising income until roughly twice the median income level and then rising sharply.

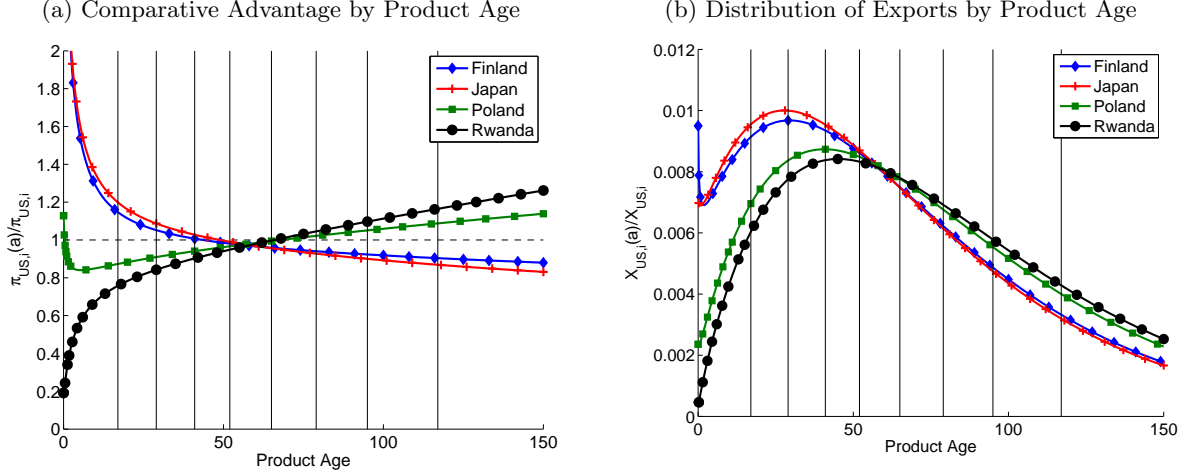
4.2.2 Product-Level Trade Patterns

Next, I demonstrate the model’s predictions for product-cycle trade patterns and evaluate the degree to which these predictions are consistent with product-level trade flow data. Figure 4 illustrates the pattern of trade that emerges in the model over the life cycle of a product for four representative countries with differing patterns of research expenditure, which are summarized in Table 3.²³ The left pane of Figure 4 depicts the evolution of countries’ comparative advantage, in expectation, as a product ages, which is measured here as the country’s expected share of US imports for products of age a , relative to the country’s share of all US imports, $\pi_{US,i}(a)/\pi_{US,i}$.²⁴ The right pane plots the expected share of each country’s total exports to the US made up of products of age a . In addition to a country’s level of comparative advantage, this measure also reflects the distribution

²³Finland and Japan devote relatively large shares of output to horizontal and vertical research, respectively. Rwanda devotes the largest share of output in the sample to learning, and Poland is an intermediate case. That the difference between the research expenditure of Finland is small is representative of the model’s prediction of a high degree of correlation in expenditure on the two forms of R&D.

²⁴The US import market is chosen for illustrative purposes because it is the largest market and has detailed product-level trade data. The pattern is similar in every destination market, though older products make up a slightly smaller share of imports in countries that perform less R&D. The share $\pi_{ni}(a) = \sum_m \eta_m \pi_{nim}(a) X_{nm}(a) / \sum_m \eta_m X_{nm}(a)$.

Figure 4: Product-Cycle Trade Patterns



of US expenditure across products differing in age.²⁵ The hump shape arises due to the offsetting forces of exponential growth of the measure of products – meaning that a given product makes up a smaller share of expenditure over time – and the accumulation of techniques over the life-cycle of a product – meaning that older products are relatively cheaper and thus attract a larger share of expenditure. The vertical lines separate the products into deciles of total U.S. imports, which aids comparison with the data below.

Three patterns emerge from this figure. First, countries which devote a higher share of GDP to R&D have an expected comparative advantage in producing newer products. Second, expenditure on a given product increases more rapidly when it is relatively young. And, third, high-R&D countries’ comparative advantage in newer products erodes slowly over time, with the expected volume of exports of a product actually increasing over a number of years and older products still making up a relatively large share of high-R&D countries’ exports. The first prediction is shared by other models that feature expanding product variety and technology diffusion, going back at least to [Krugman \(1979\)](#) and [Vernon \(1966\)](#). The second and third predictions, however, depend on countries engaging in all three types of innovative activity in the model. The second prediction arises because techniques accumulate for a given product over time. Because new techniques arrive at a constant rate, the percentage growth in expected productivity and expenditure on the product is highest when it is newest. The third prediction arises because high-R&D countries engage relatively heavily in vertical research as well as still devoting a significant share of GDP to learning, which implies that they maintain an absolute advantage over low-R&D countries for all products, and their comparative advantage falls more slowly than low-R&D country productivity rises. Further, high-R&D countries’ declining comparative advantage is offset by increasing expenditure on a product as it ages, with the latter dominating early a product’s life cycle.

²⁵Exports from i to n of age a products is given by $X_{ni}(a) = \pi_{ni}(a)X_n(a)$, where $X_n(a) = X_n \sum_m f(a, m) (\Phi_{nm}(a)/\Phi_n)^{\frac{\sigma-1}{\theta}}$.

Table 3: Example Countries

Country	s_i^h	s_i^v	s_i^ℓ
Finland	1.68	1.53	8.72
Japan	1.60	2.05	8.21
Poland	0.45	0.64	10.75
Rwanda	0.03	0.48	11.38

To test whether these patterns are present in the data, I use data on US imports of manufactured goods, categorized into 10-digit Harmonized Tariff Schedule (HTS) codes from [Schott \(2008\)](#). Unfortunately, it is not straightforward to measure the age of a product in the data, so I must attempt to do so indirectly. To do this, I take as given that the first prediction of the model holds in the data, use the revealed comparative advantage of high-R&D countries as a proxy for product age, and use this to evaluate the evidence for the second and third predictions. While this is far from the ideal test of the model’s predictions, it is reassuring that this prediction is relatively uncontroversial in models featuring product cycles and that is consistent with evidence from survey data, such as [Mansfield and Romeo \(1980\)](#) and [Mansfield et al. \(1981\)](#). Further, the latter predictions are much more relevant as they are distinct from models that do not feature all three forms of innovative activity.

To proxy for the age of a product, I calculate the following measure of the average revealed comparative advantage of low-R&D countries relative to high-R&D:

$$\text{RRCA}^b = \frac{\frac{1}{N^l} \sum_{i \in \Omega^l} X_{US,i}^b / M_{US}^b}{\frac{1}{N^h} \sum_{i \in \Omega^h} X_{US,i}^b / M_{US}^b},$$

where Ω^l and Ω^h are, respectively, the sets of low- and high-R&D countries, N^x is the number of countries in Ω^x , and $X_{US,i}^b / M_{US}^b$ is the share of US imports in HTS code b from country i .²⁶ Under the assumption that the first prediction holds, this measure is increasing, in expectation, in the age of the products that make up a product category.

I calculate this measure for all product categories in 1996 – the first year for which data on R&D expenditures is available – and sort the product categories in ascending order. To ameliorate the effects of idiosyncratic differences across products categories – and due to the fact that a product category in the data is not precisely comparable to a product in the model – after sorting the product categories by RRCA^b , I organize them into larger sets that each make up 10% of U.S. imports.²⁷

Figure 5 shows the percentage change in U.S. imports within each set of products between 1996 and 2006. The line represents the change in total US imports. The figure suggests that U.S.

²⁶High-R&D is defined as having an R&D expenditure share greater than 1.5%, and the low-R&D set is made up of the remaining countries that report R&D expenditure.

²⁷This procedure is similar to that used by [Kehoe and Ruhl \(2013\)](#) to measure changes in the extensive margin of trade flows. To ensure that each set makes up precisely 10% of US imports, the marginal product at each cut-off is split across sets. Subsequent changes in trade flows are attributed proportionally to each set.

Figure 5: Growth of US Imports by Set

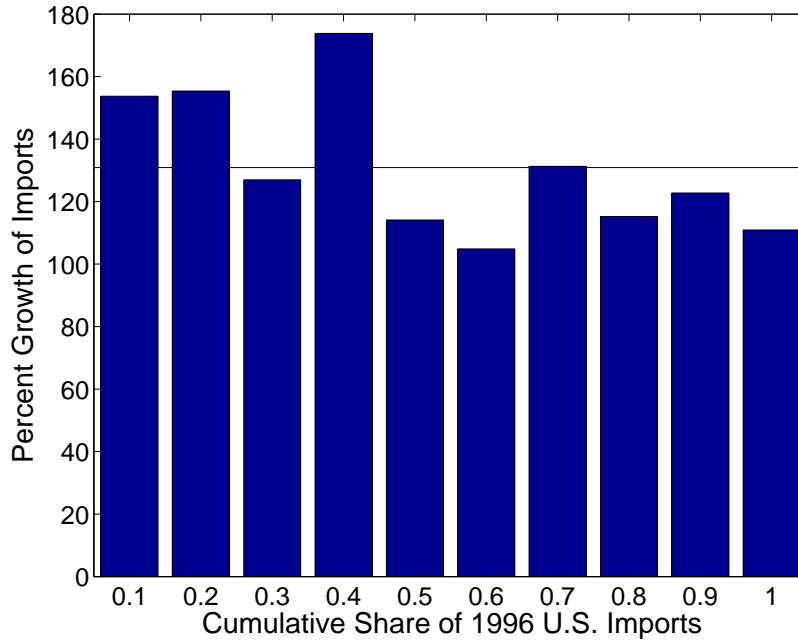
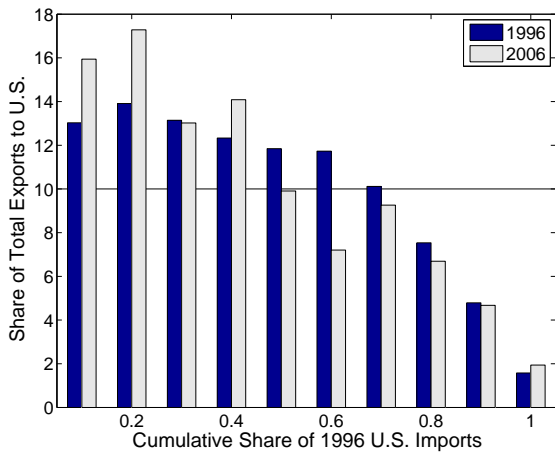
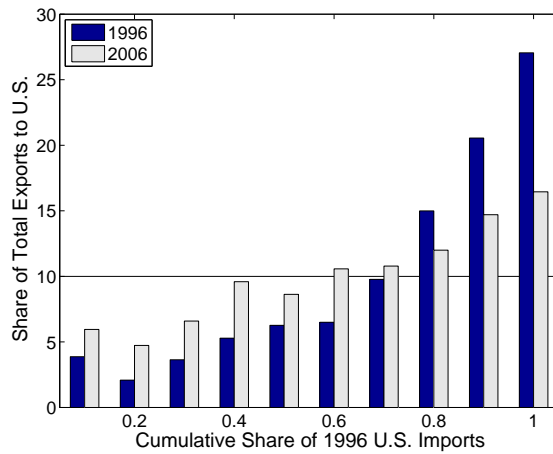


Figure 6: Growth of US Imports by Set and Country Group

(a) High R&D Countries



(b) Low R&D Countries



imports of newer products grew relatively faster than imports of older products, with imports of products in the first five sets growing by 145%, and imports of the last five growing by 117%. This is consistent with the second prediction highlighted above. Figure 6 shows the share of total exports to the U.S. made up by each set of products for both groups of countries in 1996 and 2006. This figure indicates that the exports of low-R&D countries grew much more quickly for newer products than for older products. It also indicates that exports of the newest products also grew relatively quickly for the high-R&D countries, and exports of the middle-aged products fell most rapidly as a share of total exports, consistent with Figure 4.

4.3 Knowledge Flows, Trade Flows, and the Cross-Country Income Distribution

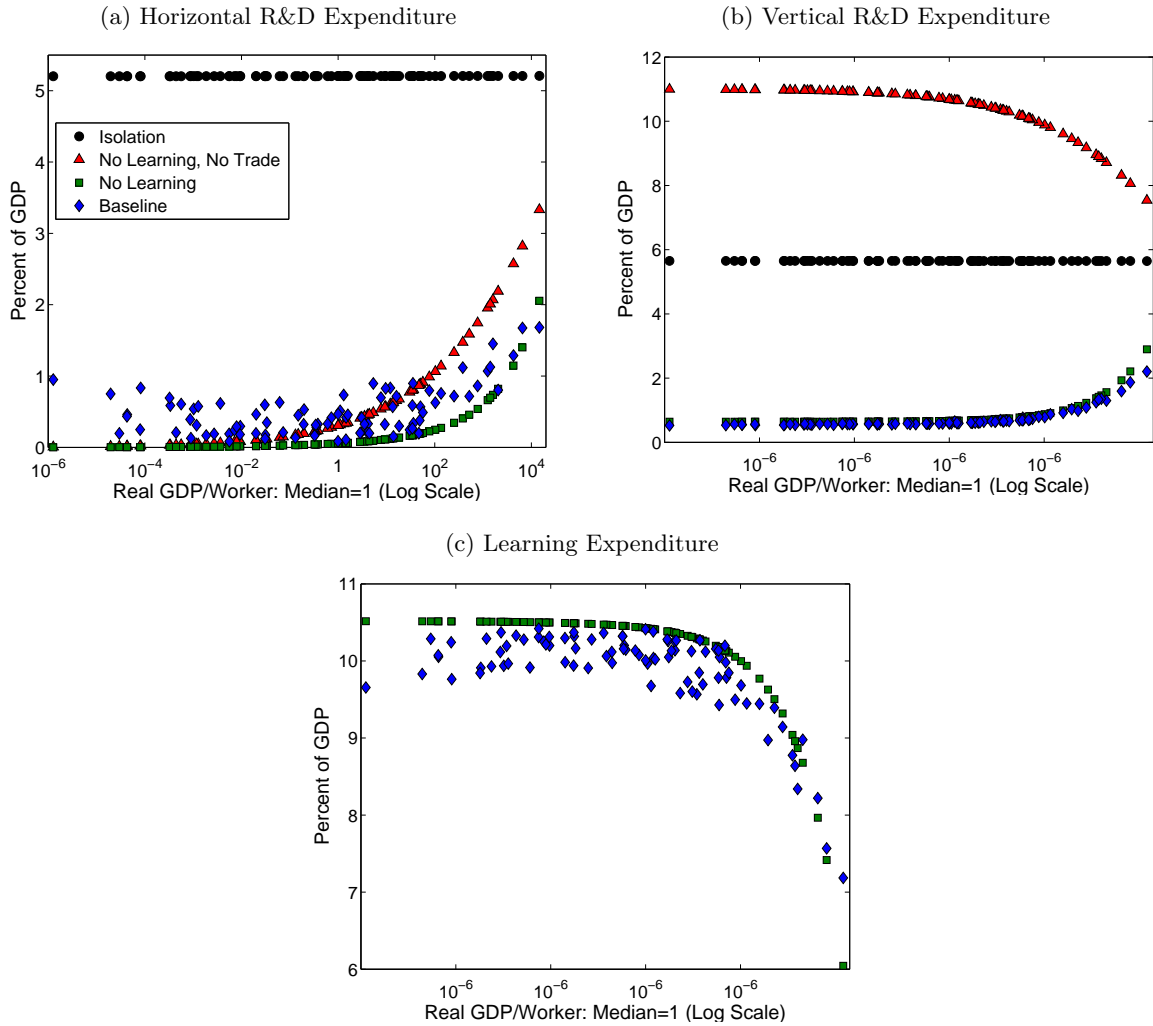
Given that the model is successful in accounting for the cross-country income distribution and patterns of R&D expenditure and product-level trade flows found in the data, I use the predictions of the model under a series of counterfactual scenarios in order to understand how the ability of both knowledge and goods to flow across national borders influence incentives to perform research and the world income distribution.

4.3.1 Knowledge and Trade Flows and Research

The first set of counterfactual experiments are designed to illustrate the effects that international trade and knowledge flows have on incentives to engage in each type of innovative activity, as discussed in Section 3.6. Figure 7 depicts the shares of output devoted to horizontal and vertical research and learning in four scenarios. In the first three, there is no international trade in goods – i.e. $d_{ni} = \infty$ for all $n \neq i$. The first is complete isolation, where countries do not know of the existence of products invented in other countries, which is equivalent to a model of N one-country worlds. In the second, the existence of new products is immediately known everywhere (as in the baseline model), but there is no learning about vertical research conducted abroad. The third is the baseline model, except without trade flows, and the fourth is the baseline model with finite trade barriers. For comparison, the figure plots research shares against real GDP per worker in isolation for all four scenarios. By definition, expenditure on learning in the first two scenarios is equal to zero.

As discussed above, every country, in isolation, chooses the same level of both horizontal and vertical research. When a country has access to the knowledge of products invented in other countries, the “smaller” is the country in terms of generation of new techniques per product, the greater is its investment in vertical research, and about half of the countries are small enough that, in the absence of trade, they devote a negligible share of output to horizontal research. The effect of the ability of researchers to learn about foreign vertical research on the incentives to engage in the other research activities is very large, especially for vertical research, where the higher returns to vertical research relative to horizontal research for countries with lower levels of potential research output in the absence of learning is reversed.

Figure 7: R&D Expenditure With and Without Trade and Knowledge Flows



In moving to the baseline case, the fall in trade barriers increases the returns to horizontal research relative to the other activities, which is clear from the increase in expenditure by almost all countries. Further, this increase is clearly heterogeneous, which indicates that the baseline case is far enough from trade autarky that relative wages are important in determining research expenditure. Countries with high research productivity and small populations increase horizontal research expenditure the most. For example, the set of countries with the largest increases in s_i^h includes Iceland, Austria, and Denmark – relatively small countries with high research productivity – as well as tiny countries such as the Maldives, Fiji, and Mauritius. The converse effect can be seen in the effect of trade on expenditure on learning, while vertical research experiences a nearly uniform fall in all countries, which reflects the fact that the returns to vertical research are not directly influenced by relative wages.

Table 4: World Income Distribution with Counterfactual Trade Costs

	$d_{ni} = \infty$ (LR)	$d_{ni} = \infty$ (SR)	Baseline —	$\min\{d_{ni}, d_{in}\}$ (SR)	$\min\{d_{ni}, d_{in}\}$ (LR)
$\% \Delta Y$	-54.6	-22.6	—	45.9	163
$\% \Delta J^{\frac{1-\gamma}{\beta(\sigma-1)}}$	-53.4	—	—	—	97.5
$\text{var}(\log(y))$	2.37	1.80	1.54	0.90	0.78
y_{90}/y_{10}	60.4	36.3	30.4	12.1	9.72

Table 5: Research Shares with Counterfactual Trade Costs

	$d_{ni} = \infty$	Baseline	$\min\{d_{ni}, d_{in}\}$
\bar{s}_i^h	0.20	0.46	0.65
\bar{s}_i^v	0.71	0.65	0.62
\bar{s}_i^l	10.98	10.75	10.59
$\text{cv}(s_i^h)$	1.73	0.84	0.54
$\text{cv}(s_i^v)$	0.58	0.52	0.44
$\text{cv}(s_i^l)$	0.071	0.064	0.055

4.3.2 Trade Barriers and Income

To understand how the world income distribution depends on trade barriers, I compare the baseline income distribution with those in two counterfactual cases, trade autarky and a case in which the asymmetric component of trade barriers has been removed. As [Waugh \(2010\)](#) argues, the latter case is likely to correspond to the case of the removal of policy-related trade barriers, while leaving in place geographic barriers, which tend to be symmetric. In both cases, I consider the “short-run” effect of a change in trade barriers – in which the stocks of products and techniques are held fixed at the levels implied by baseline balanced growth path – and a “long-run” effect, computed along the balanced growth path implied by the the new trade barriers. The former is intended to allow for comparison with other studies of the effects of changes in trade barriers in the context of static trade models.

Table 4 presents the change in total world income, the change in the component of income due to the change in the mass of products available, and two measures of the dispersion in income per worker across countries, the variance of log income per worker and the 90/10 ratio. When the state of technology in the world is held fixed, the effect of trade barriers on the distribution of income is largely in line with the predictions of static trade models, such as [Waugh \(2010\)](#). All countries benefit from lowered trade barriers, with poorer countries benefiting relatively more.²⁸ However, when research expenditure and the technological state of the world adjust to their new long-run levels, the effects of the trade barriers, both in terms of the level and dispersion of world income,

²⁸In [Waugh \(2010\)](#) the effects are quantitatively smaller, mainly because capital stocks are taken as given in his model, whereas in this framework, intermediate goods, which become cheaper as trade barriers fall, essentially play the role capital.

are amplified.

To understand these results, recall from (41) that real income per worker is a function of T_i , \tilde{T}_{ii} , π_{ii} , and J . In the short run, T_i and J are fixed. The elimination of asymmetric trade barriers benefits all countries by allowing them to specialize more fully in the set of products for which they are more productive, and it benefits poorer countries to a greater extent, since these barriers are estimated to be higher for them. Moving in the other direction toward zero trade flows, the fall in income does not covary as strongly with baseline income levels since there is not a strong relationship between π_{ii} and income in the data.

The effect of changes in \tilde{T}_{ii} turns out to be relatively small. This is for two reasons. First, in the baseline calibration, $(\sigma - 1)/\theta$ is relatively large, which implies that the additional surplus captured by producers of products that are not efficiently produced elsewhere is relatively small. Second, because trade barriers are rising or falling for all countries simultaneously, the effect on the ratio $\Phi_{nm}(a)/\Phi_n$ for all products and countries is relatively small.

In the long run, there are two additional effects. Expenditure on each type of research changes in response to the changes in trade barriers, and because tradeable goods are used in the production of research, the static gains from decreased trade barriers also imply increased output of all research activities given a set of expenditure shares. Table 5 presents the average share of income spent on each research activity as well as the coefficient of variation of each across countries. As discussed in Section 3.6, a fall in trade barriers implies that expenditure on horizontal research increases, with the relative expenditure of poorer countries increasing by more as they see a greater decrease in trade barriers. Both this and the static increase in output cause an increase in J , which further increases income proportionally for all countries. The decrease in s_i^v in most countries relative to the global increase in horizontal research causes a fall in the average number of techniques per product worldwide and implies a decrease in T_i in most countries. However, since the static output gains are greater than average for poorer countries, their gross expenditure on vertical research and learning increases, offsetting this effect. In fact, for 65% of countries, vertical research and learning expenditure increases relatively more than the global increase in the measure of products, and T_i rises. Thus, poorer countries' incomes increase by more than would be implied by the increase in J alone, while the incomes of the richest countries increase by less. The effect of a change in T_i is tempered somewhat by changes in Π_{ii} between the short run and long run, as a relative increase in T_i implies that a country becomes larger relative to the rest of the world, meaning that it has less scope for specializing in the products for which it is most efficient.

The long-run effects of the move to zero trade flows similarly amplify the static effects. The reason for the larger effect on the dispersion of income in the long run than in the short run is that the countries that are most harmed by the loss of trade flows – the smallest and least productive countries – shift to the bottom of the income distribution as a result of the static effects off the increase in trade barriers, and thus it is the poorest countries whose incomes are most affected by the long-run changes in research output.

Table 9, in Appendix A, reports the same information as Table 4 given the alternative set of

parameters based on $\theta = 6.7$, and Table 10 is the counterpart of Table 5. For changes in trade costs, the results are qualitatively similar, but the magnitudes are smaller, except for the long-run effect of a move to trade autarky. There are three reasons for this. First, a larger value of θ results in lower estimated trade barriers. Second, it implies that the effect of trade barriers on the returns to horizontal relative vertical research is greater. And, third, as is clear in (41), it dampens the effect of T_i and π_{ii} on relative incomes. In the short run, the smaller change in trade barriers and smaller effect of a change in domestic trade shares on income combine to dampen the effects of a move to symmetric trade costs. In the long run, the response of research expenditure is amplified but not by enough to cause changes in income to be greater overall. In the short run move to trade autarky, the effect of changes in domestic trade shares is mechanically lower given a higher value of θ . In the long run, the increase in vertical relative to horizontal research expenditure is much larger, so the measure of products decreases by more, and the value of T_i increases by more (or decreases by less) for all countries compared to the baseline parameterization. This benefits poorer countries relative to richer ones, leading to a smaller change in the dispersion of income across countries, but, overall, the larger change in the measure of products – whose effect on income depends on σ and not θ – dominates, so that the average country is predicted to be worse off given a higher value of θ .

4.3.3 Knowledge Flows and Income

The next set of counterfactual experiments is designed to understand how knowledge flows across countries shape the world income distribution. Table 6 depicts changes in the world income distribution under four scenarios, and Table 7 depicts the effects on research expenditure. I consider the case in which there is no learning about vertical research as well as the case in which resources invested in learning are twice as productive, which could be likened to an increase in worldwide interconnectedness due to improvements in communication and access to information via the internet. I also consider the case in which all countries are totally isolated. To distinguish the effect of isolation from knowledge of products invented abroad from that of the lack of trade flows, I include the case in which there is no learning and no trade, but entrants in each country are aware of all products that have been invented. Finally, for completeness, I include the case in which learning is twice as productive and asymmetric trade costs have been removed.

The effects of changes in the cost of learning primarily arise through changes in J and T_i , with changes in trade flows amplifying the effects somewhat. The elimination of learning increases both horizontal and vertical innovation. Thus, the measure of available products increases. This and the loss of techniques acquired through learning decreases the number of techniques available per product, especially in countries initially far from the world technological frontier. The increase in vertical research offsets this effect, but, in the case, the former effect dominates. Overall, the fall in income due to the decrease in T_i dominates the increase in income due to the increase in J , with the overall loss being much larger for countries far from the technological frontier. Thus, there is a substantial increase in the dispersion of income across countries. A halving of the cost of learning,

Table 6: World Income Distribution with Counterfactual Knowledge Flows

	Isolation ($d_{ni} = \infty$)	No Learning ($d_{ni} = \infty$)	No Learning —	Baseline —	Faster Learning —	Faster Learning $\min\{d_{ni}, d_{in}\}$
$\% \Delta Y$	-98.8	-79.2	-52.2	—	41.7	288
$\% \Delta J^{\frac{1-\gamma}{\beta(\sigma-1)}}$	-36.2	-43.5	43.5	—	-21.4	49.9
$\text{var}(\log(y))$	27.4	2.77	1.74	1.54	1.28	0.67
y_{90}/y_{10}	1.15×10^6	84.6	36	30.4	23.3	8.73

Table 7: Research Shares with Counterfactual Knowledge Flows

	Isolation ($d_{ni} = \infty$)	No Learning ($d_{ni} = \infty$)	No Learning —	Baseline —	Faster Learning —	Faster Learning $\min\{d_{ni}, d_{in}\}$
\bar{s}_i^h	5.20	0.54	3.19	0.50	0.16	0.21
\bar{s}_i^v	5.65	10.43	7.87	0.70	0.17	0.16
\bar{s}_i^l	—	—	—	9.88	10.68	10.62
$\text{cv}(s_i^h)$	—	1.25	0.44	0.68	1.69	1.20
$\text{cv}(s_i^v)$	—	0.07	0.16	0.41	1.46	1.24
$\text{cv}(s_i^l)$	—	—	—	0.059	0.046	0.038

predictably, has the opposite effect, benefitting all countries, with the gains being greater for poorer countries.

Tables 11 and 10, in Appendix A, are the respective counterparts of Tables 6 and 7 for the alternative parameterization in which $\theta = 6.7$. In predicting the effect of changes in the ease of learning on the worldwide distribution of income, θ plays a critical role. With the higher value of θ , access to foreign vertical research is not nearly as beneficial, and a decrease in the cost of learning actually leads to lower incomes in some countries, due primarily to the greater effect of worldwide learning on the returns to vertical research. When the cost of learning is halved, J falls somewhat more than in the baseline parameterization, but the effect on income is partially ameliorated by the slightly higher value of σ . The number of techniques per product in all countries increases by somewhat less, and the effect on income is further lessened due to the higher value of θ . The result is that all countries benefit less from the increased learning, and while most countries still benefit, nine of the countries closest to the technological frontier see a fall in income.²⁹ The dispersion in world income also falls by less under this parameterization.

The effect of complete isolation is predictably devastating, especially to the smallest and poorest countries. For instance, Ireland goes from being the ninth richest country in the sample to the 14th poorest, and real income per worker in the Maldives falls from 1/12th the level of the US in the baseline case to less than on billionth of this level. By contrast, US income falls by a factor of 46. When $\theta = 6.7$, average world income falls by a bit less, but isolation is even more costly for the smallest and poorest countries, increasing income inequality across countries to astronomical levels.

²⁹These countries are Finland, France, Germany, Italy, Japan, Korea, Spain, Sweden, and the United States. Finland, Germany, Italy, Japan, Korea, and Sweden also see their incomes rise with the total elimination of learning.

4.3.4 The Gains from Openness

These counterfactual experiments have assessed the effects on the world income distribution of three forms of “openness” of countries to the world economy: access to the knowledge of products invented abroad, access to vertical research aimed at developing techniques for producing these products, and access to trade in the products themselves. To assess the relative importance of each form of openness, note that, from complete isolation, allowing the first type of knowledge flows increases world income by a factor of 21. From this point, introducing access to foreign vertical research further increases world income by a factor of 2.3, while separately introducing access to international trade does so by a factor of 3.5. Introducing both together does so by a factor of 7.1. Thus, openness, broadly defined, results in a 150-fold increase in world income per worker over complete isolation, with the bulk of this gain being due to simply knowing that new products have been invented elsewhere in the world.

This also implies that access to trade and foreign vertical research are substitutes in their effects on income. For example, introducing learning in a world with no trade increases total world income by 118%, while doing so in a world with the baseline trade barriers increases world income by slightly less, 109%. Similarly, in the first case, the 90/10 ratio of income per worker falls by 29% with the introduction of learning, but it falls by only 16% with baseline trade costs. The latter result highlights the fact that it is countries far from the technological frontier that face high trade barriers that benefit the most from learning.

5 Conclusion

This paper provides a framework in which to jointly evaluate the effects of access to both foreign knowledge and foreign goods on the level and cross-country distribution of income per worker. Given its richness in being able to quantitatively address the cross-country distribution of income, aggregate and product-level bilateral trade flows, and research and development expenditure, the model provides a reasonably parsimonious and tractable framework in which to analyze the effects of international trade and knowledge flows on research incentives and the distribution of world income together in a unified framework. When parameterized to fit data on trade flows and R&D expenditure, the model’s predictions are consistent with data on income per worker, higher moments of the relationship between income and research expenditure, and product-level trade flow data; and it provides evidence that access to both international trade and knowledge flows are quantitatively important in determining the world distribution of income.

Of course, to achieve the level of richness and tractability, the model includes simplifying assumptions along a number of dimensions. These include the simple treatment of intellectual property rights and lack of an explicit channel for multinational production. As such, these remain fruitful areas of future research.

References

- Acemoglu, Daron, Philippe Aghion, and Fabrizio Zilibotti**, “Distance to the Frontier, Selection, and Economic Growth,” *Journal of the European Economic Association*, 2006, 4 (1), 37–74.
- Alvarez, Fernando and Robert E. Lucas Jr.**, “General Equilibrium Analysis of the Eaton-Kortum Model of International Trade,” *Journal of Monetary Economics*, 2007, 54 (6), 1726 – 1768.
- Anderson, James E. and Eric van Wincoop**, “Gravity with Gravitas: A Solution to the Border Puzzle,” *The American Economic Review*, 2003, 93 (1), 170–192.
- Antràs, Pol**, “Incomplete Contracts and the Product Cycle,” *American Economic Review*, 2005, 95 (4), 1054–1073.
- Arkolakis, Costas, Arnaud Costinot, and Andres Rodriguez-Clare**, “New Trade Models, Same Old Gains?,” *American Economic Review*, 2012, 102 (1), 94–130.
- Benhabib, Jess and Mark M. Spiegel**, “The Role of Human Capital in Economic Development: Evidence from Aggregate Cross-Country Data,” *Journal of Monetary Economics*, 1994, 34 (2), 143 – 173.
- Bernard, Andrew B., Jonathan Eaton, J. Bradford Jensen, and Samuel Kortum**, “Plants and Productivity in International Trade,” *The American Economic Review*, 2003, 93 (4), 1268–1290.
- Broda, Christian and David E. Weinstein**, “Globalization and the Gains from Variety,” *Quarterly Journal of Economics*, 2006, 121 (2), 541–585.
- Caselli, Francesco**, “Accounting for Cross-Country Income Differences,” *Handbook of Economic Growth*, 2005, 1, 679–741.
- Diamond, Jared M.**, *Guns, Germs, and Steel: The Fates of Human Societies*, New York: W.W. Norton, 1997.
- Dinopoulos, Elias and Peter Thompson**, “Schumpeterian Growth without Scale Effects,” *Journal of Economic Growth*, 1998, 3 (4), pp. 313–335.
- Eaton, Jonathan and Samuel Kortum**, “International Technology Diffusion: Theory and Measurement,” *International Economic Review*, 1999, 40 (3), 537–570.
- **and** – , “Technology, Trade, and Growth: A Unified Framework,” *European Economic Review*, 2001, 45 (46), 742 – 755.
- **and** – , “Technology, Geography, and Trade,” *Econometrica*, 2002, 70 (5), 1741–1779.

- **and** –, *Technology in the Global Economy: A Framework for Quantitative Analysis*, Manuscript, 2010.
- French, Scott**, “Innovation, Technology Diffusion, and Trade,” 2013, *Working Paper*.
- , “The Composition of Trade Flows and the Aggregate Effects of Trade Barriers,” 2014, *Working Paper*.
- Gollin, Douglas**, “Getting Income Shares Right,” *Journal of Political Economy*, 2002, 110 (2), pp. 458–474.
- Grossman, Gene M. and Elhanan Helpman**, “Endogenous Product Cycles,” *The Economic Journal*, 1991, 101 (408), pp. 1214–1229.
- **and** –, “Quality Ladders and Product Cycles,” *The Quarterly Journal of Economics*, 1991, 106 (2), 557–586.
- Heston, Alan, Robert Summers, and Bettina Aten**, *Penn World Table Version 7.1*, Center for International Comparisons of Production, Income and Prices at the University of Pennsylvania, November 2012.
- Howitt, Peter**, “Steady Endogenous Growth with Population and R&D Inputs Growing,” *Journal of Political Economy*, 1999, 107 (4), pp. 715–730.
- , “Endogenous Growth and Cross-Country Income Differences,” *The American Economic Review*, 2000, 90 (4), pp. 829–846.
- Jones, Charles I.**, “R&D-Based Models of Economic Growth,” *Journal of Political Economy*, 1995, 103 (4), pp. 759–784.
- Kehoe, Timothy J. and Kim J. Ruhl**, “How Important Is the New Goods Margin in International Trade?,” *Journal of Political Economy*, 2013, 121 (2), pp. 358–392.
- Kortum, Samuel S.**, “Research, Patenting, and Technological Change,” *Econometrica*, 1997, 65 (6), 1389–1419.
- Krugman, Paul R.**, “Increasing Returns, Monopolistic Competition, and International Trade,” *Journal of International Economics*, 1979, 9 (4), 469 – 479.
- Mansfield, Edwin and Anthony Romeo**, “Technology Transfer to Overseas Subsidiaries by U.S.-Based Firms,” *The Quarterly Journal of Economics*, 1980, 95 (4), 737–750.
- , **Mark Schwartz, and Samuel Wagner**, “Imitation Costs and Patents: An Empirical Study,” *Economic Journal*, December 1981, 91 (364), 907–18.
- Mayer, Thierry and Soledad Zignago**, “Notes on CEPII’s Distances Measures: The GeoDist Database,” Working Papers 2011-25, CEPII 2011.

- Mehra, Rajnish and Edward C. Prescott**, “The Equity Premium: A Puzzle,” *Journal of Monetary Economics*, 1985, 15 (2), 145 – 161.
- Nelson, Richard R.**, “The Role of Knowledge in R&D Efficiency,” *The Quarterly Journal of Economics*, 1982, 97 (3), pp. 453–470.
- Parente, Stephen L. and Edward C. Prescott**, *Barriers to Riches*, Cambridge: MIT Press, 2002.
- Peretto, Pietro F.**, “Technological Change and Population Growth,” *Journal of Economic Growth*, 1998, 3 (4), 283–311.
- Pierce, Justin R. and Peter K. Schott**, “Concording U.S. Harmonized System Categories Over Time,” Working Paper 14837, National Bureau of Economic Research April 2009.
- Rivera-Batiz, Luis A. and Paul M. Romer**, “Economic Integration and Endogenous Growth,” *The Quarterly Journal of Economics*, 1991, 106 (2), pp. 531–555.
- Samuelson, Paul A.**, “The Transfer Problem and Transport Costs, II: Analysis of Effects of Trade Impediments,” *The Economic Journal*, 1954, 64 (254), pp. 264–289.
- Schott, Peter K.**, “The Relative Sophistication of Chinese Exports,” *Economic Policy*, 2008, 23 (53), 5–49.
- Segerstrom, Paul S.**, “Endogenous Growth without Scale Effects,” *The American Economic Review*, 1998, 88 (5), pp. 1290–1310.
- Silva, J. M. C. Santos and Silvana Tenreyro**, “The Log of Gravity,” *The Review of Economics and Statistics*, 2006, 88 (4), pp. 641–658.
- Vandenbussche, Jérôme, Philippe Aghion, and Costas Meghir**, “Growth, Distance to Frontier and Composition of Human Capital,” *Journal of Economic Growth*, 2006, 11 (2), 97–127.
- Vernon, Raymond**, “International Investment and International Trade in the Product Cycle,” *The Quarterly Journal of Economics*, 1966, 80 (2), 190–207.
- Waugh, Michael E.**, “International Trade and Income Differences,” *American Economic Review*, 2010, 100 (5), 2093–2124.
- and **Ina Simonovska**, “The Elasticity of Trade: Estimates and Evidence,” 2013, *Working Paper*.
- Young, Alwyn**, “Growth without Scale Effects,” *Journal of Political Economy*, 1998, 106 (1), pp. 41–63.

Appendices

A Additional Tables

Table 8: Parameter Values: $\theta = 6.7$

Symbol	Value	Target
r	0.07	Rate of return on equity
g_L	0.015	World population growth rate
γ	0.33	Ratio of labor compensation to gross output in non-tradable sectors
β	0.22	Ratio of labor compensation to gross output in manufacturing
σ	4.8	GDP per capita growth rate
θ	6.7	Trade cost elasticity
λ	0.73	Average R&D expenditure
α^ℓ	2.7	Average R&D expenditure among G5 countries
α^h	0.73	Relative number of firms engaged in product innovation

Table 9: World Income Distribution with Counterfactual Trade Costs ($\theta = 6.7$)

	$d_{ni} = \infty$ (LR)	$d_{ni} = \infty$ (SR)	Baseline —	$\min\{d_{ni}, d_{in}\}$ (SR)	$\min\{d_{ni}, d_{in}\}$ (LR)
$\% \Delta Y$	-61.4	-16.1	—	27.9	111
$\% \Delta J^{\frac{1-\gamma}{\beta(\sigma-1)}}$	-63.4	—	—	—	79.8
$\text{var}(\log(y))$	1.72	1.60	1.48	0.95	0.86
y_{90}/y_{10}	34.6	32.5	28.7	13.6	11.1

Table 10: Research Shares with Counterfactual Trade Costs ($\theta = 6.7$)

	$d_{ni} = \infty$	Baseline	$\min\{d_{ni}, d_{in}\}$
\bar{s}_i^h	0.07	0.44	0.67
\bar{s}_i^v	1.12	0.72	0.60
\bar{s}_i^l	9.61	9.47	9.32
$\text{cv}(s_i^h)$	3.51	1.18	1.06
$\text{cv}(s_i^v)$	0.39	0.50	0.48
$\text{cv}(s_i^l)$	0.079	0.078	0.082

Table 11: World Income Distribution with Counterfactual Knowledge Flows ($\theta = 6.7$)

	Isolation ($d_{ni} = \infty$)	No Learning ($d_{ni} = \infty$)	No Learning —	Baseline —	Faster Learning —	Faster Learning $\min\{d_{ni}, d_{in}\}$
$\% \Delta Y$	-98.7	-72.4	-12.0	—	5.49	117
$\% \Delta J^{\frac{1-\gamma}{\beta(\sigma-1)}}$	-46.3	-59.4	48.0	—	-26.6	22.0
$\text{var}(\log(y))$	62.5	1.85	1.59	1.48	1.33	0.78
y_{90}/y_{10}	1.96×10^9	39.3	31.8	28.7	24.2	9.57

Table 12: Research Shares with Counterfactual Knowledge Flows ($\theta = 6.7$)

	Isolation ($d_{ni} = \infty$)	No Learning ($d_{ni} = \infty$)	No Learning —	Baseline —	Faster Learning —	Faster Learning $\min\{d_{ni}, d_{in}\}$
\bar{s}_i^h	4.48	0.113	1.81	0.439	0.135	0.191
\bar{s}_i^v	5.58	10.3	8.71	0.721	0.131	0.115
\bar{s}_i^l	—	—	—	9.47	10.2	10.1
$\text{cv}(s_i^h)$	—	3.08	1.02	1.18	2.46	2.05
$\text{cv}(s_i^v)$	—	0.04	0.18	0.50	2.14	1.95
$\text{cv}(s_i^l)$	—	—	—	0.078	0.057	0.05

Table 13: Country-Level Parameter Values

	L_i	ex_i	$\left(\frac{\alpha_i Y_i}{\alpha_{US} Y_{US}}\right)^\lambda$	L_i	ex_i	$\left(\frac{\alpha_i Y_i}{\alpha_{US} Y_{US}}\right)^\lambda$	L_i	ex_i	$\left(\frac{\alpha_i Y_i}{\alpha_{US} Y_{US}}\right)^\lambda$		
Albania	1,357	-1.30	0.005	Gambia	679	-4.00	0.003	Portugal	5,599	0.68	0.177
Argentina	18,234	0.92	0.163	Germany	41,418	2.77	0.833	Rep. of Korea	23,892	1.93	0.853
Australia	10,715	2.05	0.401	Ghana	9,386	-0.88	0.009	Romania	10,500	0.45	0.074
Austria	4,086	0.98	0.241	Greece	4,678	0.02	0.181	Russian Fed.	73,948	1.53	0.209
Bangladesh	68,596	-0.09	0.019	Hungary	4,237	1.06	0.138	Rwanda	4,837	-3.80	0.005
Bel., Lux.	4,939	2.49	0.162	Iceland	172	-0.86	0.071	Saint Lucia	79	-1.72	0.001
Benin	3,224	-2.81	0.007	India	448,571	1.00	0.083	Samoa	70	-3.06	0.011
Bolivia	4,097	-1.79	0.024	Indon., Mal., Sing.	125,516	5.39	0.026	Sao Tome, Princ.	55	-3.67	0.001
Bosnia, Herz.	1,651	0.16	0.007	Ireland	2,166	4.10	0.039	Senegal	4,731	-1.13	0.009
Botswana	934	-0.54	0.006	Israel	2,764	0.78	0.143	Serbia	3,885	-1.04	0.064
Brazil	97,396	1.41	0.235	Italy	25,016	1.66	1.013	Slovakia	2,679	2.38	0.017
Bulgaria	3,358	-0.32	0.058	Jamaica	1,274	-0.51	0.010	Slovenia	1,025	-0.35	0.116
Burundi	4,164	-3.35	0.003	Japan	66,437	2.67	1.376	Solomon Islds	202	-5.16	0.003
Côte d'Ivoire	7,506	-1.09	0.027	Jordan	1,425	-1.34	0.047	South Africa	17,666	1.52	0.177
Cambodia	7,139	-1.11	0.011	Kenya	13,411	-0.55	0.014	Spain	21,763	1.73	0.551
Cameroon	7,227	-2.57	0.039	Lebanon	1,268	-1.58	0.053	Sri Lanka	8,824	-0.71	0.060
Canada	17,621	1.23	0.393	Macedonia	895	-1.96	0.025	Sudan	12,310	-0.89	0.012
Central America	15,407	0.49	0.055	Malawi	6,214	-2.18	0.003	Swaziland	465	-2.30	0.053
China, H.K., Mac.	775,248	3.72	0.257	Maldives	154	-1.11	0.003	Sweden	4,638	1.34	0.398
Colombia	18,920	-0.19	0.073	Mali	3,421	-1.62	0.002	Tanzania	19,170	-0.66	0.002
Croatia	2,008	-1.08	0.108	Mauritius	557	0.15	0.018	Thailand	36,510	2.79	0.063
Cuba	4,834	-1.64	0.059	Mexico	44,296	1.91	0.072	Tunisia	3,514	-0.50	0.052
Czech Rep.	5,184	0.40	0.255	Morocco	10,878	-0.64	0.041	Turkey	26,854	0.82	0.223
Denmark	2,913	0.88	0.167	Namibia	785	-0.74	0.031	USA	152,766	5.16	1.000
Dominican Rep.	4,064	-1.83	0.101	Netherlands	8,712	3.62	0.068	Uganda	11,741	-1.56	0.005
Ecuador	6,745	-0.41	0.033	New Zealand	2,178	1.42	0.188	United Kingdom	31,139	2.22	0.468
Ethiopia	36,647	-0.98	0.002	Pakistan	57,557	-0.09	0.030	Uruguay	1,606	-0.88	0.056
Fiji	354	-0.63	0.012	Paraguay	2,748	-0.63	0.014	Viet Nam	48,597	1.01	0.031
Finland	2,652	0.54	0.483	Peru	13,231	0.26	0.075	Zambia	5,174	-0.07	0.005
France	29,257	1.85	0.715	Poland	17,172	0.47	0.237	Zimbabwe	5,992	-0.01	0.002

Notes: Labor force is expressed in thousands of persons. The values reported under the ex_i heading are equal to $-\theta(ex_i + \frac{1}{N} \sum_i ex_i)$ and are comparable to the distance coefficients reported in table 2. "Central America" includes Belize, Costa Rica, Guatemala, Honduras, El Salvador, and Panama. The term in the final column represents a country's capacity to produce research, relative to that of the United States.

B Data

B.1 GDP and Labor Force Statistics

Data on real GDP and the size of the labor force are from version 7.1 of the Penn World Tables (Heston et al. (2012)). The measure of real GDP used is total PPP Converted GDP, based on the Geary-Khamis method, at current prices in 2006. The size of the labor force for each country is computed, as in Caselli (2005), as $RGDPCH*POP/RGDPWOK$ in 2006, where $RGDPCH$ is PPP converted GDP per capita, computed using the chain method, at 2005 constant prices; $RGDPWOK$ is PPP converted GDP per worker in the same units; and POP is population.

B.2 Labor Compensation Shares

Labor compensation as a share of output in manufacturing and non-tradeable sectors is calculated from data obtained from the STAN database available from the OECD for 2006.

B.3 Research & Development Statistics

Data on national expenditure on R&D as a percentage of GDP are taken from the World Development Indicators (WDI) database available from the World Bank for 2006. Data on the number of enterprises engaged in horizontal and vertical research is from the 2006 Community Innovation Survey available from Eurostat. The survey asked enterprises whether they introduce “new or significantly improved goods” or services and whether they introduced “new or significantly improved methods of manufacturing or producing goods or services”, “new or significantly improved logistics, delivery or distribution methods for you inputs, goods or services”, or “new or significantly improved supporting activities for your processes” between 2004 and 2006. Data is available for 21 European countries which are also in the sample of countries for which the model is computed.

B.4 Manufacturing Output

Data on total gross manufacturing output is necessary to estimate border costs and the country-specific research productivity terms. The data is obtained from two sources. Where it is available, the data is taken from the Industrial Statistics Database (INDSTAT4), 2011 Edition, CD-ROM available from the United Nations Industrial Development Organization. Where data for 2006 is not available but is available for other years, the 2006 value is interpolated using the predicted values of a regression of log output on country and year fixed effects. Where no data is available for a country, gross manufacturing output is imputed from total manufacturing value added obtained from WDI. Manufacturing value added is scaled up by a factor of 3.19 based on a cross-sectional regression of gross output on value added with no constant term.

B.5 Trade Data

Product-level, bilateral trade data, which is used to estimate bilateral trade costs, is taken from the U.N. Comtrade database classified into six-digit Harmonized System (HS) product codes. For 2006, the database contains trade flow data for 146 countries that report data using the HS2002 classification and 19 countries that report using the HS1996 classification. HS1996 codes are converted to HS2002 codes using the correspondence table available from the U.N. Statistics Division.³⁰ These 165 reporting countries report trade with an additional 77 non-reporting countries and territories. However, to ensure a complete trade flow matrix, only reporting countries are included in the sample.

For pairs of reporting countries, bilateral trade flows are typically reported in both directions by both countries. Trade flows reported by the importing country were used except for country pairs in which the importer reported no trade flows across all product categories, but trade flows were reported by the exporter. Such trade flows make up 1.4% of the observations. This results in a dataset of 165 countries, 5,222 product codes, and 6,854,937 non-zero bilateral, product-level trade flow observations.

Data on product-level U.S. imports are from Schott (2008). I use trade data from the years 1996 – the first year for which R&D data are available – and 2006 – the year for which the remaining data used to parameterize the model are taken. Because U.S. HTS codes change often, in order to compare product-level trade over time, I use the concordance developed by Pierce and Schott (2009). This data is merged with data on R&D expenditure from WDI. For countries in the sample that do not report R&D data until later years, the first available observation is used. Cuba and Zambia are omitted because they do not export to the U.S. in all years during this period. The remaining sample of imports from countries that report R&D expenditure is made up of 52 countries and 13,622 HTS codes or sets of codes concorded between 1996 and 2006.

The bilateral relationship variables used to estimate trade costs are from the Gravity dataset available from CEPII (see Mayer and Zignago (2011)). The variables used in the estimation are population weighted distance (*distw*), whether countries share a border (*contig*), whether they share a common official language (*comlang_off*), whether they have ever had a colonial link (*colony*), and whether they are currently members of a common regional trade agreement (*rta*).

B.6 Constructing the Sample

In the model, a country is defined by the size of its labor force, L_i , its average productivity in research, α_i , and the bilateral trade barriers it faces, d_{ni} . Thus, for a country to be included in the sample, data must be available for labor force size, manufacturing output, and product-level trade flows, with the latter two being used to estimate trade barriers and research productivity.

To combine the trade flow data with manufacturing output data, trade in non-manufacturing HS codes are dropped from the dataset. These are identified using the mapping from HS2002 codes

³⁰This is available for free download from the following url: <http://unstats.un.org/unsd/cr/registry/regdntransfer.asp?f=232>.

to ISIC (revision 3) codes available from the U.N. Statistics Division.³¹ All HS codes not mapped to ISIC 2-digit industries 15-37 are dropped. This reduces the number of HS codes in the sample to 4,702 and the number of observations to 6,500,799. Domestic absorption of domestic manufacturing output, X_{ii} , is then calculated as manufacturing output minus total manufacturing exports to all countries (including non-reporters).

Lack of manufacturing output data for all countries reduces the sample size to 149 countries. Lack of data from the Penn World Tables further reduces the sample size to 130 countries. To avoid problems related to entrepot trade, several sets of countries are grouped into regional blocs. These include, Belgium and Luxembourg; China, Hong Kong, and Macao; Singapore, Malaysia, and Indonesia; and Belize, Costa Rica, Guatemala, Honduras, El Salvador, and Panama. Where there was no natural regional bloc to join, countries for which exports in excess of manufacturing output was reported were dropped from the sample.³² These two steps together reduced the sample to 109 countries and regional blocs.

Finally, since the model is not well-equipped to explain the level of income per worker for countries which obtain a large fraction of GDP from natural resource extraction, I drop countries from the sample if non-manufacturing industrial value added accounts for more than 30% of GDP. The following countries are dropped based on this criterion: Algeria, Bhutan, Brunei Darussalam, Chile, Gabon, Guinea, Iran, Mauritania, Mongolia, Nigeria, Norway, Saudi Arabia, Trinidad and Tobago, United Arab Emirates, Venezuela, and Yemen. Such data is missing for Botswana, Greece, Israel, Oman, Qatar, Romania, and Syria. Of these countries, Oman, Qatar, and Syria are dropped because a large share of these countries' exports is made up of petroleum products. The remaining countries are retained. These steps result in a final sample of 90 countries and 3,812,144 positive product-level bilateral trade flows (37,663,020 including zero trade flows). This sample accounts for 93% of world GDP recorded by the Penn World Tables in 2006.

C Computation

Solving the model numerically requires evaluating the integrals given in (35) and (39). While these integrals cannot be solved analytically, they can be computed numerically using standard techniques. Given the form of the integrand, Gauss-Laguerre quadrature was found to perform quite well. The double integrals defined by (39) can be converted to single integrals using the change of variable $z = s + a$ and integrating a out analytically.

Given a full set of parameter values, I use the following algorithm to solve for the balanced growth path equilibrium:

1. Set initial values for s_i^x , for all i and x .
2. Set initial values for wages, w_i , for all i .

³¹This is available for free download from the following url: <http://unstats.un.org/unsd/cr/registry/regdntransfer.asp?f=184>.

³²These included The Bahamas, Cyprus, Estonia, Guyana, Madagascar, Malta, Mozambique, Philippines, St. Vincent, Vanuatu.

3. Set initial values for Φ_n , for all n .
4. Using (33), (34), and (35), update the values of Φ_n .
5. Repeat 4 until the values of Φ_n converge.
6. Calculate π_{ni} according to (25).
7. Using (31), update w_i , i.e.

$$w'_i = w_i + \delta \left(\sum_n \pi_{ni} w_n L_n - w_i L_i \right), \quad \delta \in (0, 1).$$

8. Repeat 4 - 7 until the values of w_i converge.
9. Calculate V_i^x , for all i and x , according to (39).
10. Using (40), update the values of s_i , i.e.

$$s'_i = \frac{1 - \gamma}{1 + \beta\theta} \frac{g_Y}{R^h} \alpha_i^x (s_i^x)^\lambda \left(\frac{Y_i}{Y} \right)^\lambda V_i^x$$

11. Repeat 4 - 10 until the values of s_i converge.

D Proofs

D.1 Productivity Distribution

Here, I show that if (1) a given technique has productivity Z drawn from a Pareto distribution given by

$$\Pr(Z > z) = (z/\underline{z})^{-\theta}, \quad z \in [\underline{z}, \infty),$$

and (2) the number of techniques available is distributed Poisson with parameter $\underline{z}^{-\theta} T_i^j$, then the distribution of the productivity of the best and second best techniques is given by (20). This proof closely follows Eaton and Kortum (2010).

Denote the productivity of the highest productivity technique Z_1 and that of the second highest productivity technique Z_2 . If there are k techniques, the probability that $Z_1 \leq z_1$ and $Z_2 \leq z_2$, where $\underline{z} \leq z_2 \leq z_1$, is equal to the probability that *any* one of the k techniques has a productivity between z_2 and z_1 *and* all of the other techniques have productivity less than z_2 plus the probability that all of the techniques have productivity less than z_2 . Thus, for $k \geq 1$,

$$\begin{aligned} F(z_1, z_2|k) &= \Pr(Z_1 \leq z_1, Z_2 \leq z_2|k) \\ &= \Pr(z_2 \leq Z_1 \leq z_1, Z_2 \leq z_2|k) + \Pr(Z_1 \leq z_2, Z_2 \leq z_2|k) \\ &= k \left(z_2^{-\theta} \underline{z}^\theta - z_1^{-\theta} \underline{z}^\theta \right) \left(1 - z_2^{-\theta} \underline{z}^\theta \right)^{k-1} + \left(1 - z_2^{-\theta} \underline{z}^\theta \right)^k, \end{aligned}$$

and for $k = 0$, $F(z_1, z_2|0) = 1$.

Since k is drawn from a Poisson distribution, this probability, unconditional on k , is

$$\begin{aligned}
F(z_1, z_2) &= \sum_{k=0}^{\infty} \frac{e^{-T_i^j \underline{z}^{-\theta}} (T_i^j \underline{z}^{-\theta})^k}{k!} F(z_1, z_2|k) \\
&= \sum_{k=1}^{\infty} \frac{e^{-T_i^j \underline{z}^{-\theta}} (T_i^j \underline{z}^{-\theta})^k}{k!} k \left(z_2^{-\theta} \underline{z}^{\theta} - z_1^{-\theta} \underline{z}^{\theta} \right) \left(1 - z_2^{-\theta} \underline{z}^{\theta} \right)^{k-1} \\
&\quad + \sum_{k=0}^{\infty} \frac{e^{-T_i^j \underline{z}^{-\theta}} (T_i^j \underline{z}^{-\theta})^k}{k!} \left(1 - z_2^{-\theta} \underline{z}^{\theta} \right)^k \\
&= T_i^j \left(z_2^{-\theta} - z_1^{-\theta} \right) e^{-T_i^j \underline{z}^{-\theta}} \sum_{k'=0}^{\infty} \frac{\left(T_i^j \underline{z}^{-\theta} \left(1 - z_2^{-\theta} \underline{z}^{\theta} \right) \right)^{k'}}{k'!} \\
&\quad + e^{-T_i^j \underline{z}^{-\theta}} \sum_{k=0}^{\infty} \frac{\left(T_i^j \underline{z}^{-\theta} \left(1 - z_2^{-\theta} \underline{z}^{\theta} \right) \right)^k}{k!} \\
&= T_i^j \left(z_2^{-\theta} - z_1^{-\theta} \right) e^{-T_i^j z_2^{-\theta}} + e^{-T_i^j z_2^{-\theta}},
\end{aligned}$$

which is equivalent to (20). While the term \underline{z} has disappeared entirely, the support of $F(z_1, z_2)$ still only includes values greater than \underline{z} . However, it is straightforward to show that defining $\alpha_i^v = \underline{z}^{-\theta} \tilde{\alpha}_i^v$ and $\alpha_i^\ell = \underline{z}^{-\theta} \tilde{\alpha}_i^\ell$, for all i , implies a value of T_i^j that is proportional to $\underline{z}^{-\theta}$, as is assumed above. Then, letting $\underline{z} \rightarrow 0$ allows the productivity of techniques to take on any value in $(0, \infty)$, which is mathematically convenient for what follows.

Intuitively, as $\underline{z} \rightarrow 0$, it becomes infinitely easy to obtain a technique, but techniques are almost surely of zero productivity. However, the probability of obtaining a technique with productivity above a certain threshold is independent of \underline{z} . Since low-productivity techniques are never used in equilibrium, those below an arbitrarily low threshold can be ignored, allowing the minimum productivity of techniques to be safely sent to zero.

D.2 Price and Trade Flow Results

In this section, I prove the results listed in Section 2.7. These proofs draw heavily on those in Bernard et al. (2003) and Eaton and Kortum (2010).

D.2.1 Result 1

Using (9) and (20), the lowest and second lowest costs of delivering a product from i to n are distributed according to the following:

$$\begin{aligned}
\tilde{G}_{ni}^j(c_1, c_2) &= \Pr\left(C_{ni}^{j1} > c_1, C_{ni}^{j2} > c_2\right) \\
&= \Pr\left(Z_i^{j1} < c_i d_{ni}/c_1, Z_i^{j2} < c_i d_{ni}/c_2\right) \\
&= F_i^j(c_i d_{ni}/c_1, c_i d_{ni}/c_2) \\
&= \left[1 + T_i^j(c_i d_{ni})^{-\theta}(c_2^\theta - c_1^\theta)\right] e^{-T_i^j(c_i d_{ni})^{-\theta} c_2^\theta},
\end{aligned}$$

where $0 \leq c_1 \leq c_2$. The probability that the lowest cost of delivering a product to n from *anywhere in the world* is greater than c_1 and that the second lowest cost is greater than c_2 is equal to the probability that the lowest cost from any one country is between c_1 and c_2 and the lowest cost from every other country is greater than c_2 plus the probability that the lowest cost from every country is greater than c_2 . Thus,

$$\begin{aligned}
\tilde{G}_n^j(c_1, c_2) &= \Pr\left(\min_i \{C_{ni}^{j1}\} > c_1, \min_{i \neq i^*} \{C_{ni^*}^{j2}, \min\{C_{ni}^{j1}\}\} > c_2\right) \\
&= \sum_{i=1}^N \left(\tilde{G}_{ni}^j(c_1, c_2) - \tilde{G}_{ni}^j(c_2, c_2)\right) \prod_{m \neq i} \tilde{G}_{nm}^j(c_2, c_2) + \prod_{i=1}^N \tilde{G}_{ni}^j(c_2, c_2) \\
&= \sum_{i=1}^N T_i^j(c_i d_{ni})^{-\theta}(c_2^\theta - c_1^\theta) e^{-T_i^j(c_i d_{ni})^{-\theta} c_2^\theta} \prod_{m \neq i} e^{-T_i^j(c_m d_{nm})^{-\theta} c_2^\theta} + \prod_i e^{-T_i^j(c_i d_{ni})^{-\theta} c_2^\theta} \\
&= \left[1 + \Phi_n^j(c_2^\theta - c_1^\theta)\right] e^{-\Phi_n^j c_2^\theta}.
\end{aligned}$$

D.2.2 Result 2

The lowest cost of delivering a product to n from i , unconditional on the value of the second lowest cost, is distributed according to

$$\tilde{G}_{ni}^{j1}(c_1) = \tilde{G}_{ni}^j(c_1, c_1) = e^{-T_i^j(c_i d_{ni})^{-\theta} c_1^\theta}$$

The probability that the lowest cost is that of a producer from i is obtained by integrating over all values of C_n^j for which i is the low-cost provider:

$$\pi_{ni}^j = - \int_0^\infty \prod_{m \neq i} \tilde{G}_{nm}^{j1}(c) d\tilde{G}_{ni}^{j1} = \frac{T_i^j(c_i d_{ni})^{-\theta}}{\Phi_n^j},$$

where the negative accounts for the fact that \tilde{G}_{ni}^{j1} is a complementary CDF.

D.2.3 Result 3

Though tedious, it is straightforward to show that

$$\begin{aligned} \Pr(C_n^{j1} > c_1, C_n^{j2} > c_2 | C_n^{j1} = C_{ni}^{j1}) \\ &= \frac{1}{\pi_{ni}^j} \left[\int_{c_2}^{\infty} \prod_{m \neq i} \tilde{G}_{nm}^{j1}(c) d\tilde{G}_{ni}^{j1}(c) + (\tilde{G}_{ni}^j(c_1, c_2) - \tilde{G}_{ni}^j(c_2, c_2)) \prod_{m \neq i} \tilde{G}_{nm}^j(c_2, c_2) \right] \\ &= \tilde{G}_n^j(c_1, c_2). \end{aligned}$$

D.2.4 Result 4

The joint pdf of C_n^{j1} and C_n^{j1} is

$$g_n^j(c_1, c_2) = \frac{\partial^2 \tilde{G}_n^j(c_1, c_2)}{\partial c_1 \partial c_2} = \theta^2 (\Phi_n^j)^2 (c_1 c_2)^{\theta-1} e^{-\Phi_n^j c_2^\theta}.$$

The distribution of the ratio C_n^{j2}/C_n^{j1} , conditional on $C_n^{j2} = c_2$, for $0 \leq C_n^{j1} \leq c_2$ is given by

$$\begin{aligned} \Pr(C_n^{j2}/C_n^{j1} < m | C_n^{j2} = c_2) &= \Pr(c_2/m < C_n^{j1} < c_2 | C_n^{j2} = c_2) \\ &= \frac{\int_{c_2/m}^{c_2} g_n^j(c_1, c_2) dc_1}{\int_0^{c_2} g_n^j(c_1, c_2) dc_1} \\ &= \frac{c_2^\theta - (c_2/m)^\theta}{c_2^\theta} \\ &= 1 - m^{-\theta}. \end{aligned}$$

Since the distribution, conditional on $C_n^{j2} = c_2$, is independent of c_2 , the unconditional distribution is identical.

D.2.5 Result 5

First, it is useful to derive an expression for the expectation, $E[(p_n^j)^{1-\sigma}]$. From (10),

$$p_n^j = C_n^{j2} \cdot \min \{1, \bar{m} C_n^{j1}/C_n^{j2}\}.$$

Using this along with the distribution of C_n^{j1}/C_n^{j2} , derived above, the expectation can be expressed as

$$\begin{aligned} E[(p_n^j)^{1-\sigma}] &= E[(C_n^{j2})^{1-\sigma}] (1 - \bar{m}^{-\theta}) + \int_{\bar{m}}^{\infty} E[(\bar{m} C_n^{j2}/m)^{1-\sigma}] \theta m^{-\theta-1} dm \\ &= E[(C_n^{j2})^{1-\sigma}] \left((1 - \bar{m}^{-\theta}) + \frac{\theta}{1 + \theta - \sigma} \bar{m}^{-\theta} \right), \end{aligned}$$

where the last step uses the fact that the distributions of C_n^{j2} and C_n^{j1}/C_n^{j2} are independent.

Using the expression for the joint distribution of C_n^{j1} and C_n^{j2} derived above, the unconditional distribution of C_n^{j2} can be derived:

$$\begin{aligned} g_n^{j2}(c_2) &= \int_0^{c_2} g_n^j(c, c_2) dc \\ &= \theta (\Phi_n^j)^2 c_2^{2\theta-1} e^{-\Phi_n^j c_2^\theta}. \end{aligned}$$

Using this, $E[(C_n^{j2})^{1-\sigma}]$ can be expressed as

$$\begin{aligned} E[(C_n^{j2})^{1-\sigma}] &= \int_0^\infty c_2^{1-\sigma} g_n^{j2}(c_2) dc_2 \\ &= \int_0^\infty \theta (\Phi_n^j)^2 c_2^{\sigma+2\theta} e^{-\Phi_n^j c_2^\theta} dc_2 \\ &= (\Phi_n^j)^{\frac{\sigma-1}{\theta}} \int_0^\infty x^{\frac{1-\sigma+2\theta}{\theta}-1} e^{-x} dx \\ &= (\Phi_n^j)^{\frac{\sigma-1}{\theta}} \Gamma\left(\frac{1-\sigma+2\theta}{\theta} - 1\right), \end{aligned}$$

where $\Gamma(\cdot)$ is the gamma function. Thus,

$$\begin{aligned} E[(p_n^j)^{1-\sigma}] &= \left(1 + \frac{\sigma-1}{1+\theta-\sigma} \bar{m}^{-\theta}\right) \Gamma\left(\frac{1-\sigma+2\theta}{\theta} - 1\right) (\Phi_n^j)^{\frac{\sigma-1}{\theta}} \\ &\equiv \hat{\gamma}^{1-\sigma} (\Phi_n^j)^{\frac{\sigma-1}{\theta}} \end{aligned}$$

Define the pdf for the price of product j in n as $f_n^j(p)$. It is useful to relabel a product by its realized price in n rather than its original label j . In this case, (6) can be expressed as the integral overall all possible prices, weighted by the measure of products which have each realized price:

$$\begin{aligned} P_n^{1-\sigma} &= \int_0^\infty p^{1-\sigma} \int_0^J f_n^j(p) dj dp \\ &= \int_0^J \int_0^\infty p^{1-\sigma} f_n^j(p) dp dj \\ &= \int_0^J E[(p_n^j)^{1-\sigma}] dj \\ &= \hat{\gamma}^{1-\sigma} \int_0^J (\Phi_n^j)^{\frac{\sigma-1}{\theta}} dj \\ &\equiv \hat{\gamma}^{1-\sigma} \Phi_n^{\frac{\sigma-1}{\theta}}. \end{aligned}$$

D.2.6 Result 6

Using (5) and the expressions above for the values of $E \left[(p_n^j)^{1-\sigma} \right]$ and $P_n^{1-\sigma}$, the expected value of expenditure by n on a product is given by

$$\begin{aligned} E [X_n^j] &= E \left[\left(\frac{p_n^j}{P_n} \right)^{1-\sigma} X_n \right] \\ &= E \left[(p_n^j)^{1-\sigma} \right] \frac{X_n}{P_n^{1-\sigma}} \\ &= \left(\frac{\Phi_n^j}{\Phi_n} \right)^{\frac{\sigma-1}{\theta}} X_n. \end{aligned}$$

D.2.7 Result 7

The expected value of profits for the low-cost provider of a product to n is given by

$$\begin{aligned} E [\Pi_n^j] &= E \left[\left(1 - \frac{C_n^{j1}}{p_n^j} \right) \left(\frac{p_n^j}{P_n} \right)^{1-\sigma} X_n \right] \\ &= E [X_n^j] - E \left[\frac{C_n^{j1}}{p_n^j} (p_n^j)^{1-\sigma} \right] \frac{X_n}{P_n^{1-\sigma}}. \end{aligned}$$

Following closely the proof of result 5,

$$\begin{aligned} E \left[\frac{C_n^{j1}}{p_n^j} (p_n^j)^{1-\sigma} \right] &= \int_1^{\bar{m}} E [m^{-1} (C_n^{j2})^{1-\sigma}] \theta m^{-\theta-1} dm + \bar{m}^{-1} \int_{\bar{m}}^{\infty} E [(\bar{m} C_n^{j2} / m)^{1-\sigma}] \theta m^{-\theta-1} dm \\ &= E [(C_n^{j2})^{1-\sigma}] \left(\frac{\theta}{1+\theta} (1 - \bar{m}^{-\theta-1}) + \frac{\theta}{1+\theta-\sigma} \bar{m}^{-\theta-1} \right) \\ &= \frac{\theta}{1+\theta} E [(C_n^{j2})^{1-\sigma}] \left(1 + \frac{\sigma-1}{1+\theta-\sigma} \bar{m}^{-\theta} \right). \end{aligned}$$

Thus,

$$\begin{aligned} E [\Pi_n^j] &= E [X_n^j] - \frac{\theta}{1+\theta} E [X_n^j] \\ &= \frac{1}{1+\theta} E [X_n^j]. \end{aligned}$$