Technology, Trade, and Growth: A Unified Framework

Jonathan Eaton and Samuel Kortum*
Boston University and NBER

January 2001

Abstract

Our recent work examines the links among innovation, technology, trade, and growth. One strand focuses on research activity, technology diffusion, and growth. The other examines technology and trade. In this paper we exploit the common treatment of technology in these two strands to provide a parsimonious model of innovation, growth, and trade. We examine the effect of lower geographic barriers to trade on research and the effects of scale and research productivity on relative incomes.

Key words: trade, growth, technology, research

JEL classification: F11; F17; O31; O33; O40

*Department of Economics, Boston University, 270 Bay State Road, Boston, MA 02215, USA. E-mail: jeaton@bu.edu; telephone: (617) 353-4142; fax: (617) 353-4449. We received useful comments on an earlier draft from Kei-Mu Yi. We gratefully acknowledge the support of the National Science Foundation under Grant No. SBR 9601765.

1 Introduction

We report on a line of research examining the links among innovation, technology, trade, and growth. The research has been aimed at developing a theoretical structure that suits itself to empirical implementation. It goes on to quantify this structure using aggregate and microlevel data.

One strand of this work focuses on the links among research activity, patenting, technology diffusion, and growth. Kortum (1997) develops a single-country search-theoretic model of innovation and growth to explain puzzling trends in productivity, patents, and R&D activity in the United States. Eaton and Kortum (1999) extend the model to a multi-country world with international technology diffusion. They measure the extent of technology diffusion among the five major research economies by fitting the model to data on research, productivity, and bilateral patent applications.[1]

Another strand of research examines the links between technology and trade. Eaton and Kortum (2000a) develop a model of bilateral trade which they fit to trade and price data among the OECD. They explore the extent to which the benefits of technology are shared through the exchange of products. Bernard, Eaton, Jensen, and Kortum (2000) go on to analyze the relationship between the productivity of individual producers and their ability to penetrate export markets. They use the model to reconcile bilateral trade data with facts about the export behavior of individual U.S. plants.

Common to all of these papers is a treatment of technology that allows a unified theory of innovation, trade, and growth. In each of them we augmented the basic model to capture features of the data crucial to the application at hand. These additions were not central to the fundamental workings of the models, however, and may have even obscured their underlying unity.

[1]Eaton and Kortum (1997) use the model to explain the post-World War II productivity convergence of these five countries.

Here we present a parsimonious framework that encompasses the basic structure in each of these papers. We do so in order to reveal the connections between the forces driving innovation and productivity, on the one hand, and the implications of technology for trade, on the other. Hence we focus on the theory rather than our quantitative findings.[2]

One new issue we confront is the effect of increased openness on research incentives and growth. The model incorporates two offsetting effects on research incentives. One is the potentially much larger market that any successful innovator can exploit through exports. The other is the greater difficulty in coming up with an idea that not only advances the domestic state of the art but also competes successfully with technologies available through imports. The first effect encourages research while the second discourages it.

Our model provides a baseline in which these two forces exactly cancel: Steady-state research intensities are invariant to geographic barriers to trade. While there are static gains from trade, there are no dynamic gains through the accumulation of technology. While generalizations of the model can destroy this stark neutrality, there is no presumption about the market-enlarging effects of trade on innovation.

A second issue we address is the role of pure scale and research productivity in determining living standards. In general, a country's real wage depends not only on how good its workers are at coming up with ideas, but on how many workers it has to come up with them. In the extreme case of autarky, relative real wages depend on relative labor forces weighted by research productivity. Lowering geographic barriers to trade, however, benefits countries with smaller labor forces disproportionately, so that countries with more labor start to lose their edge. Moving to the opposite extreme of no barriers to trade, a more populous country's

[2]While we do not exploit it in this paper, a major feature of our framework is its ability to capture the heterogeneity behind aggregate data. In the case of inventions, some ideas seem such breakthroughs they are patented widely; less spectacular ideas are patented only at home; most ideas enter the dustbin. In the case of efficiency, some producers are so efficient they are able to sell widely; less efficient ones sell only at home; the really inefficient are driven out of business. Hence our framework can, for example, link data on aggregate productivity to international patent counts and aggregate trade flows to data on the export behavior of individual plants.

advantage from having more people to generate ideas is exactly offset by its relatively lower gains from trade. Hence, with frictionless trade, relative wages depend only on relative research productivities. Relative labor forces do not matter.

Section 2 below presents the static trade model. Section 3 embeds it into a dynamic model of research and growth. Section 4 offers some concluding remarks.

2 Trade

We consider a world with a unit continuum of consumption goods and N countries. We first describe the technology for production and trade, and then turn to preferences and the gains from trade.

2.1 The Technology Frontier

Any country is capable of producing any good, although it might not be very good at it. The efficiency with which country i produces good j is $z_i(j)$, where $j \in [0,1]$ and $i \in \{1,...,N\}$. We assume labor is the only input to production, so the z's represent labor productivities.[3] We refer to $\{z_i(j)|j \in [0,1]\}$ as country i's technological frontier since it represents the best techniques of production available there.

All the agents in the economy know the technological frontier. But rather than keeping track of all of the $z_i(j)$'s ourselves we treat them as realizations of random variables Z_i drawn from a distribution F_i. We can thus represent a country's technological frontier with only the small number of parameters of the distribution F_i.

It turns out that a particular functional form for F_i yields a very tractable formulation of trade among countries at any moment. In particular, we assume that F_i has the Fréchet form:

$$F_i(z) = \Pr[Z_i \leq z] = e^{-T_i z^{-\theta}}, \quad z \geq 0. \tag{1}$$

[3] In Eaton and Kortum (2000a), production combines labor and intermediate inputs so the z's represent total factor productivities. In either case $z_i(j)$ can also be interpreted as the quality of good j produced in country i.

The distribution has two parameters, each of which is central to the analysis.[4]

The accumulated technology parameter T_i, which we treat as country-specific, governs the location of the distribution. Countries with higher levels of T tend to be more efficient at producing any good j. As we show in the next section, T_i measures the history of inventions that have been absorbed into country i's technology. In a trade context T reflects absolute advantage.

The parameter θ governs the amount of variation in efficiency around a country's mean, with lower values of θ implying more heterogeneity. In a trade context θ reflects the scope for exploiting comparative advantage.[5]

We make the Ricardian assumption that labor is perfectly mobile across activities and regions within a country but does not move across countries. Before turning to international trade, consider what the model implies within a country that constitutes a single labor and goods market, but with technology differing among its K regions.

Suppose each region k within the country has a technology frontier F_k with parameters T_k and θ. With full mobility of goods and labor, any good j will be produced only in the region with highest efficiency. Hence the country's efficiency level is $z(j) = \max\{z_1(j), ..., z_K(j)\}$. Since each $z_k(j)$ is the realization of a random variable drawn from F_k, $z(j)$ itself is the realization of a random variable Z with distribution:

$$F(z) = \Pr[Z \leq z] = \prod_{k=1}^{K} \Pr[Z_k \leq z] = \prod_{k=1}^{K} F_k(z) = \prod_{k=1}^{K} e^{-T_k z^{-\theta}} = e^{-T z^{-\theta}}, \qquad (2)$$

where $T = \sum_{k=1}^{K} T_k$. Hence we can summarize the country's overall accumulated technology T as the sum of the accumulated technology T_k in each region.

Our assumption on the form of the distribution of the technology frontier is, of course, very special. As we showed above, however, it does embody the fundamental property of replicating

[4]In its basic assumptions our model resembles Dornbusch, Fischer, and Samuelson's (1977) two-country Ricardian model of trade with a continuum of goods. By specifying the technology frontier as Frechet we have made it a useful tool for the empirical analysis of trade patterns in a world of many countries.

[5]To see what the parameters imply for the location and variation, note that $\ln Z_i$ has mean $\ln \gamma + (1/\theta) \ln T_i$ and coefficient of variation $\pi/(\theta\sqrt{6})$, where γ is Euler's constant.

itself under aggregation across space. As we show in the next section, under certain conditions it also replicates itself under aggregation over time. As a result, it can be derived from a dynamic process of innovation.

2.2 The Pattern of Trade

In the case of international trade, labor in different countries may earn different wages and moving goods between countries can be costly. We denote the wage in country i as w_i (using w_N as numeraire) and represent the cost of moving goods from country i to country n by the parameter d_{ni}. Here d_{ni} reflects geographic barriers to the movement of goods of the standard iceberg variety: Delivering one unit of a good to country n from country i requires shipping $d_{ni} \geq 1$ from country i. (For each i we set $d_{ii} = 1$.) We interpret these costs broadly to include not only tariffs and transport costs but also additional impediments to the free flow and use of foreign goods such as costs of search, negotiation, and adoption.

The cost of buying good j from country i in country n is:

$$c_{ni}(j) = \frac{w_i d_{ni}}{z_i(j)}.$$

Our assumption about the $z_i(j)$'s means that $c_{ni}(j)$'s are realizations drawn from:

$$G_{ni}(c) = 1 - e^{-T_i (w_i d_{ni})^{-\theta} c^{\theta}}.$$

Note that the distribution of these costs takes into account not only country $i's$ accumulated technology, but also the cost of paying labor there and delivering goods to country n.

We assume that country n buys each good only from the cheapest source (as is the case under perfect or Bertrand competition). The cost of good j in country i is then $c_n(j) = \min\{c_{n1}(j), ..., c_{nN}(j)\}$. By the same logic used in (2), the lowest costs available in country n, the $c_n(j)$'s, are realizations from:

$$G_n(c) = 1 - e^{-\Phi_n c^{\theta}} \tag{3}$$

where $\Phi_n = \sum_{i=1}^{N} T_i(w_i d_{ni})^{-\theta}$, each country's accumulated technology downweighted by its labor cost and the cost of delivery to n. The cost parameter Φ_n reflects country n's ability to exploit technology around the world through international trade.

Remarkably, as we show in Eaton and Kortum (2000a), the distribution (3) applies not only to the cost of any good that country n buys unconditional upon source, it applies conditioning on source as well. That is, given that country i crosses the hurdle of selling in country n, its cost there has the same distribution as any other country actually selling there. A country that is more backward, remote, or high-wage will simply cross the hurdle less often.

The probability π_{ni} that country i is the cheapest source of a particular good in country n is simply i's share of Φ_n:

$$\pi_{ni} = \frac{T_i(w_i d_{ni})^{-\theta}}{\sum_{k=1}^{N} T_k(w_k d_{nk})^{-\theta}}. \tag{4}$$

Since there are a continuum of goods, this probability also represents the fraction of goods that country n buys from i.

Except in the zero-probability event of a tie, country n buys each good from only one source. However, in the presence of geographic barriers ($d_{ni} > 1$), more than one country may produce the same good.

Under perfect competition, of course, prices correspond to costs. Since consumers in country n will then face the same price distribution for any good they actually buy regardless of where it comes from, π_{ni} represents the fraction of spending in country n devoted to goods from country i.[6]

Expenditure shares reflect the range of goods that n buys from i. By allowing for endogenous specialization our Ricardian model contrasts with the monopolistic competition and Armington approaches popular in general equilibrium trade modeling. Under the Armington approach each country's goods are distinct simply by assumption. There is no sense, then,

[6]As we show in Bernard, Eaton, Jensen, and Kortum (2000), this feature of the model generalizes very handily to Bertrand competition.

in which countries compete with each other to supply a particular good. Under monopolistic competition each country endogenously specializes in a different set of commodities, so again there is no head-to-head competition good by good. In either of these other approaches, geographic barriers influence how much of a good is sold, but not whether the good is sold at all.

2.3 The Gains from Trade

To derive the model's implications for trade and welfare, we need to specify preferences. We make the simplest assumption that utility of a representative consumer in country i is a Cobb-Douglas function across the consumption of individual varieties. Hence:

$$U = \exp \int_0^1 \ln x_i(j) dj$$

where $x_i(j)$ is the consumption of good j in country i.[7] Since expenditure shares are independent of price, under this assumption the π_{ni} correspond to trade shares under any form of competition in which only the cheapest source sells.

We can now demonstrate how trade, by allowing countries to exploit each other's technologies, confers gains. The exact price index in country n, P_n, under Cobb-Douglas and perfect competition, is:

$$P_n = \exp \int_0^1 \ln c_n(j) dj = \exp \int_0^\infty \ln(c) dG_n(c) = \gamma \Phi_n^{-1/\theta}, \tag{5}$$

where, again, γ is Euler's constant.[8]

Dividing the wage by this price index, invoking (4), gives the real wage:

$$\frac{w_n}{P_n} = \gamma^{-1} \left[\sum_{i=1}^N \frac{T_i}{T_n} \left(\frac{w_i}{w_n} \right)^{-\theta} d_{ni}^{-\theta} \right]^{1/\theta} T_n^{1/\theta} = \gamma^{-1} \left(\frac{T_n}{\pi_{nn}} \right)^{1/\theta}. \tag{6}$$

[7] The analysis generalizes very readily to constant elasticity of substitution preferences, as we show in Eaton and Kortum (2000a).

[8] As we show in Eaton and Kortum (2000a), under CES preferences the price index is identical except for the definition of γ, which remains independent of n. As we show in Bernard, Eaton, Jensen, and Kortum (2000) under Bertrand competition the price index is also identical, except for a further (also n invariant) redefinition of γ. Hence under either generalization the price index in country n remains proportional to $\Phi_n^{-1/\theta}$.

Under autarky the real wage in country n is just $\gamma^{-1}T_n^{1/\theta}$, so only domestic technology matters. Accessing foreign technology through trade augments domestic technology by the factor $1/\pi_{nn} = 1/(1-I_n)$, where I_n is n's import share. Hence, the gains from trade are an increasing function of the share of purchases from abroad, and knowing θ, these gains can be quantified.

It turns out that the real wage corresponds to the average efficiency of active producers in a country. Equation (6) thus illustrates the role of trade in pushing workers into activities in which the country is more efficient, while driving the least efficient producers out of business. Aggregate productivity thus rises as countries open their borders to imports.

While we can infer a country's trade gains from its import share, this share itself depends not only on the parameters governing technology and geographic barriers, but also on endogenously determined wages. For these, we turn to the conditions for labor-market equilibrium in each country.

With a common profit share across countries (which, under perfect competition, is zero and which, in the case of Bertrand competition that we explore below, turns out to be $1/(1+\theta)$), the condition for labor market equilibrium in country i is:

$$w_i L_i^P = \sum_{n=1}^{N} \pi_{ni} w_n L_n^P, \tag{7}$$

where L_i^P are the number of production workers in country i.

Equations (4), (6), and (7) together determine trade shares, wages, and welfare as functions of the parameters of the technology distributions T_i and θ, geographic barriers d_{ni}, and labor supplies L_i^P.

As we show in Eaton and Kortum (2000a), the case of zero gravity (all $d_{ni} = 1$) yields a simple closed form solution of the complete model. The relative wage in country i is:

$$\frac{w_i}{w_N} = \left(\frac{T_i/L_i^P}{T_N/L_N^P}\right)^{1/(1+\theta)}. \tag{8}$$

With all $d_{ni} = 1$ the price level is the same everywhere, so (8) also corresponds to the relative real wage. In contrast, from (6), the relative real wage under autarky is:

$$\frac{w_i/P_i}{w_N/P_N} = \left(\frac{T_i}{T_N}\right)^{1/\theta}. \tag{9}$$

Hence under autarky it is the absolute level of technology that determines a country's relative welfare while under free trade it is how much technology it has per worker. (Thus small countries gain relatively more from trade than large ones.)

3 Trade and Growth

So far we have analyzed trade given the technology frontiers (1). We now show how these frontiers emerge from an underlying process of innovation and how the incentives to innovate divide the workforce between producers and researchers. In making the model dynamic we allow technologies T to evolve over time. (We treat geographic barriers d_{ni} and the parameter θ as fixed over time.)

Researchers draw ideas about how to produce goods. A given researcher in country i draws ideas at a Poisson rate α_i, a parameter representing research productivity. An idea is the realization of two random variables. One is the good j to which it applies, which is drawn from the uniform distribution over $[0, 1]$. The other is the efficiency $q(j)$ with which it enables good j to be produced, which is drawn from the Pareto distribution $H(q) = 1 - q^{-\theta}$. The efficient technology $z_i(j)$ for producing good j in country i is the best idea for producing it yet discovered. A new idea is never adopted unless it surpasses the current state of the art $z_i(j)$. Even if it does, it may not be able to survive competition from abroad.

In the absence of technology diffusion the number of ideas that a country has available to it depends on the history of research effort there. If the measure of researchers in country i at time s is R_{is} the total stock of ideas there is $T_{it} = \alpha_i \int_0^t R_{is} ds$. Because we have a unit interval of varieties, the number of ideas for producing a specific good is distributed Poisson with parameter T_{it}. As we now show, this T turns out to be the same one that enters the technology frontier (1).

The Poisson arrival of ideas implies that the probability of k ideas for producing a particular good by date t in country i is $(T_{it})^k e^{-T_{it}}/k!$. If there have been k ideas, the probability that the best one is below z is $[H(z)]^k$. Summing over all possible numbers of ideas:

$$F_{it}(z) = \sum_{k=0}^{\infty} \frac{{T_{it}}^k e^{-T_{it}}}{k!} H(z)^k = e^{-T_{it}[1-H(z)]} = e^{-T_{it}z^{-\theta}}, \quad z \geq 1.$$

3.1 Innovation

Although a researcher in country i gets ideas as a Poisson process with parameter α_i, most of those ideas will have a quality below the frontier technology. Even many which surpass the local frontier may not hold up to competition from abroad. The probability that an idea of quality q will be competitive in country n, i.e. that $w_{it}d_{ni}/q$ is the lowest cost in country n, is given by $1 - G_{nt}(w_{it}d_{ni}/q)$, where G is given by (3). More generally, the probability that the idea will undercut the lowest cost by a factor $m \geq 1$ is $1 - G_{nt}(mw_{it}d_{ni}/q)$. Integrating over the Pareto distribution of idea quality:

$$b_{nit}(m) = \int_1^{\infty} [1 - G_{nt}(mw_{it}d_{ni}/q)]\, dH(q) \approx \frac{1}{\Phi_{nt}(mw_{it}d_{ni})^{\theta}}. \tag{10}$$

Of course setting $m = 1$ gives the probability of entering the market at all.[9]

Incorporating the trade share expression (4) from the previous section into this expression implies that the probability of an idea from i having a market in n is given by:

$$b_{nit}(1) = \frac{\pi_{nit}}{T_{it}}. \tag{11}$$

For a country i invention to succeed in market n it must cross two hurdles. First, the idea has to surpass the previous state of the art in country i, which it does with probability $1/T_{it}$. Conditional on being frontier technology in country i, the idea has to beat out the foreign competition in market n, which it does with probability π_{nit}.

[9] The approximation in (10) has to do with how we handle efficiency levels below 1. Note that our derivation of $F_{it}(z)$ is only defined for $z \geq 1$ while (1) is defined for all $z \geq 0$. Kortum (1997) shows that this problem can be safely ignored. The reason is that $F_{it}(1)$ approaches zero as T_{it} gets large.

Note that as time passes and the stock of existing ideas in country i grows, it gets harder and harder to come with new ideas that are better than the best existing ones.

3.2 Markups

How are researchers compensated for their efforts? Under perfect competition producers would charge marginal cost, so there would be no return to innovating. Following the quality ladders model (Grossman and Helpman, 1991; Aghion and Howitt, 1991) we make the Bertrand assumption that the lowest cost producer of each good claims the entire market for that good, charging the highest markup that keeps any competitor at bay. In our model the markup of any successful entrant or surviving incumbent is a random variable M with Pareto distribution H, as we now show.

¿From (10), the probability of selling in market n at all is $b_{nit}(1)$ while the probability of selling in that market with markup higher than $m \geq 1$ is $b_{nit}(m)$. Thus, conditional on selling at all, the distribution of the markup is

$$\Pr[M \leq m | M \geq 1] = \frac{b_{nit}(1) - b_{nit}(m)}{b_{nit}(1)} = H(m),$$

which is simply Pareto.

The markup distribution does not vary with time, destination, or source. It follows that the distribution of the markup at any date $s \geq t$ for ideas that are still competing at that time is also Pareto. Imagine a cohort of ideas discovered at date t in country i which have a market in country n. As time proceeds, some of the ideas will be surpassed by subsequent ideas so will drop out of market n. The probability of remaining in that market by time $s \geq t$ is b_{nis}/b_{nit}. (From now on we use b to denote $b(1)$, the probability of remaining in the market at all.) The ideas remaining in the market may have their markups whittled down over time by new ideas. Amazingly, the selection process, which favors ideas with larger markups, exactly offsets the whittling down effect. The distribution of the markup for market survivors remains exactly the same over time.

3.3 Profits

Let Y_{nt} denote total expenditure at date t. Since preferences are Cobb-Douglas, Y_{nt} is also the rate of expenditure per variety in country n. Since all firms selling in market n charge a markup drawn from H, total profits earned by either domestic or foreign firms selling there are:

$$\Pi_{nt} = Y_{nt} \int_0^1 \left[1 - m(j)^{-1}\right] dj = Y_{nt} \int_1^\infty (1 - m^{-1}) dH(m) = \frac{Y_{nt}}{1+\theta}.$$

What is the profit earned by a firm producing in country n? Its cut of market k profits is simply its trade share there π_{knt}. Hence its total earnings around the world are:

$$\sum_{k=1}^{N} \pi_{knt} \Pi_{kt} = \sum_{k=1}^{N} \frac{\pi_{knt} Y_{kt}}{1+\theta} = \frac{Y_{nt}}{1+\theta} \tag{12}$$

where the last equality follows from the balanced trade condition that $Y_{nt} = \sum_{k=1}^{N} \pi_{knt} Y_{kt}$. Hence Π_{nt} also corresponds to profits earned by firms in country n. The profit share in the economy is consequently $1/(1+\theta)$. Since this share is constant across countries, our labor market equilibrium condition (7) stands.

3.4 Research Incentives

Putting together the pieces from the previous section, the expected present discounted value of an idea from country i that succeeds in country n at date t is:

$$V_{nit} = P_{it} \int_t^\infty e^{-\rho(s-t)} \frac{\Pi_{ns}}{P_{is}} \frac{b_{nis}}{b_{nit}} ds, \tag{13}$$

where ρ is the discount factor (assumed constant). Profits are discounted by the discount rate but also by the dwindling probability of remaining in business in the market. (Here P_{is} is the price index in country i at time s.)

But the probability of a researcher in country i finding an idea that succeeds in breaking into market n at time t is just b_{nit}. Summing across all markets, the expected value of a discovery in country i is simply:

$$V_{it} = \sum_{n=1}^{N} b_{nit} V_{nit} = P_{it} \int_t^{\infty} e^{-\rho(s-t)} \sum_{n=1}^{N} \frac{\Pi_{ns}}{P_{is}} b_{nis} ds. = \frac{P_{it}}{1+\theta} \int_t^{\infty} e^{-\rho(s-t)} \frac{Y_{is}}{P_{is}} \frac{1}{T_{is}} ds, \qquad (14)$$

where the last equality follows from the condition for balanced trade. Note that, even though our model incorporates international trade, we can ascertain the value of an idea originating in country i by looking only at country i's own income and accumulated technology.

As in the quality ladders model, in equilibrium workers divide themselves between research and production. An equilibrium with workers engaged in both activities requires that the wage of a production worker w_{it} equal the expected return to research, $\alpha_i V_{it}$.

3.5 Steady State Growth

For simplicity we consider a steady state in which a constant share $r_i = R_{it}/L_{it}$ of the labor force in country i engages in research. To allow such a steady state to emerge we assume that labor forces everywhere grow at a constant rate g_L.

With constant r_i the level of technology in country i evolves according to:

$$\dot{T}_{it} = \alpha_i r_i L_{it}$$

so that T/L converges to $t_i = \alpha_i r_i / g_L$. Hence in steady state T_i grows at rate g_L in each country i.

We can rewrite the labor market equilibrium condition (7) as:

$$w_{it}(1 - r_i) = \sum_{n=1}^{N} \frac{t_i (w_{it} d_{ni})^{-\theta}}{\sum_{k=1}^{N} t_k (w_{kt} d_{nk})^{-\theta} (L_{kt}/L_{nt})} w_{nt}(1 - r_n).$$

Since L_{kt}/L_{nt} is constant over time, the solution for relative wages is time invariant. Since w_N is numeraire, wages themselves are time invariant.

Increased real income comes about from falling prices. Differentiating the price index (5) with respect to time, prices everywhere fall at rate

$$-\frac{\dot{P}_i}{P_i} = \frac{1}{\theta}\frac{\dot{T}}{T} = \frac{g_L}{\theta}.$$

Substituting the results for steady state into the expression for the value of an idea in country i (14) we get:

$$V_{it} = \frac{g_L}{\alpha_i r_i}\frac{(1-r_i)w_{it}}{\theta\rho - g_L}. \tag{15}$$

Here we see that we need to impose the condition $\rho > g_L/\theta$ to obtain a finite value.

Equating the value of doing research in country i, $\alpha_i V_{it}$ with the wage w_{it} there we get:

$$r_i = r = \frac{g_L}{\theta\rho}. \tag{16}$$

This expression for equilibrium research intensity is identical to the one chosen by a social planner (see Kortum, 1997).

Not surprisingly, research intensity is greater the higher the rate of labor force growth and the lower the discount factor. Moreover research intensity is greater the smaller is θ. (With smaller θ, successful inventions on average constitute larger advances over the preceding state of the art.)

More surprising is that research intensity does not depend on country size, research productivity, or openness. While access to foreign markets increases the potential profits that a successful idea can earn, competition from foreign inventions makes it more difficult to have a marketable idea in the first place.[10]

Although a large country has the same research intensity, the absolute number of researchers matters for the research stock. In steady state the relative technology levels are:

$$\frac{T_i}{T_N} = \frac{\alpha_i L_i}{\alpha_N L_N}. \tag{17}$$

[10] Rivera-Batiz and Romer (1991) also find invariance of research intensity to the opening up of trade (but not ideas) in their "knowledge-driven" specification. Their analysis is based on the monopolistic competition model of trade. More critically, their model, unlike ours, delivers scale effects on growth.

A country's accumulated technology depends on the size of its labor force weighted by research productivity. The scale of a country may therefore affect its level of welfare.

To see how trade influences the outcome, we can combine (17) with our expressions for relative real wages under autarky (9) and under zero gravity (8). With autarky scale does matter: A country's relative real wage depends on its research productivity multiplied by the size of its labor force. But countries with smaller labor forces gain more from trade. With the complete elimination of trade barriers the greater gains from trade of smaller countries exactly offsets their scale disadvantage in producing ideas. With zero gravity, relative real wages depend only on relative research productivity, with scale playing no role.

4 Conclusion

The framework summarized above is readily extended to examine a much wider range of issues. For example, incorporating technology diffusion would allow innovations to be exploited abroad not only through exports but through the movement of ideas themselves. This extension introduces a number of challenging problems which we begin to explore in Eaton and Kortum (1997, 1999). Exports of capital goods embodying technological advances are another conduit for spreading the benefits of technology, as we explore in Eaton and Kortum (2000b).

These extensions address only a small fraction of the myriad issues relating to technology that emerge in the global economy. Our approach delivers a tractable structure for attacking them.

References

Aghion, P. and P. Howitt, 1992, A model of growth through creative destruction, *Econometrica* 60, 323-351.

Bernard, A., J. Eaton, J. B. Jensen, and S. Kortum, 2000, Plants and productivity in international trade, *NBER Working Paper* 7688.

Dornbusch, R., S. Fischer, and P.A. Samuelson, 1977, Comparative advantage, trade, and payments in a ricardian model with a continuum of goods, *American Economic Review* 67, 823-839.

Eaton, J. and S. Kortum, 1997, Engines of growth: Domestic and foreign sources of innovation, *Japan and the World Economy* 9, 235-259.

Eaton, J. and S. Kortum, 1999, International technology diffusion: Theory and measurement, *International Economic Review* 40, 537-570.

Eaton, J. and S. Kortum, 2000a, Technology, geography, and trade, http://econ.bu.edu/eaton.

Eaton, J. and S. Kortum 2000b, Trade in capital goods, *NBER Working Paper* 8070, forthcoming in *European Economic Review.*

Grossman, Gene M. and Elhanan Helpman, 1991, Quality ladders in the theory of growth, *Review of Economic Studies* 58, 43-61.

Kortum, S., 1997, Research, patenting and technological change, *Econometrica* 65, 1389-1419.

Rivera-Batiz, L. A. and P. Romer, 1991, Economic integration and endogenous growth, *Quarterly Journal of Economics* 107, 531-556.