14.662 Spring 2015, Lecture Note 5: Ricardian Models of Trade

David Autor

MIT and NBER

March 7, 2015

1 Introduction

Why and what do countries trade, and why are international trade flows changing? How can we distinguish labor market impacts of trade from technology–and how do these two interact? How can we identify the effects of trade on labor markets? What's the evidence? Those are the questions we'll try to address in this section of the class.

I've taught the Hecksher-Ohlin model of trade in graduate labor classes since 1999 to familiarize students with how trade and labor markets may interact in theory and in practice. Hecksher-Ohlin models are driven by differences in factor intensities across countries. Exogenous supply differences—often skill supplies—determine comparative advantage among countries. Several years ago, I also began teaching Ricardian models. These models are driven by differences in productivity across countries. Comparative advantage comes from technology differences (as well as geography). Ricardian models have so come to dominate contemporary trade theory and trade empirics that I plan to shelve teaching of Hecksher-Ohlin models in 14.662. I will, make available the notes on this subject for self-study (if interested). If I'm underestimating the residual demand for these lectures, please let me know.

2 Ricardian models of trade

There are a number of key empirical regularities in trade that are basically incomprehensible within the canonical Hecksher-Ohlin model, in which the sole determinant of trading patterns is differences in factor abundance. These facts are:

1. Trade between countries diminishes with distance

2. Large countries trade *less* relative to GDP, but trade relatively more in absolute terms. Thus, country size is quite important.

3. Countries import more from larger countries

4. Prices vary across locations, with greater price differences between countries that are further apart

5. Productivities within the same industry appear to differ across countries—suggesting that trade/specialization is at least partly based on technology differences, not simply factor endowments. (If we don't attribute this excess specialization to intrinsic productivity diffs, we need to invoke something like increasing returns—which is how Krugman earned his Nobel.)

6. Factor rewards *do not* appear to be equalized across countries (that is, FPE does not hold)

How does the Ricardian model differ from the canonical Hecksher-Ohlin model?

1. Many goods, many countries (as opposed to two-by-two)

2. Productivity varies across industries within countries (rather than common technologies)

3. Countries also differ in their overall productivity (so there, is a country's TFP as well as its industry specific productivity)

4. Trade is costly (in HO models, it's common to assume away transport costs)

What we hope to get out

1. Understanding of how productivity differences affect trade

2. An understanding of how productivity growth in one country affects labor markets in others

3. Rationalization of some basic facts: distance, size

4. Ability to use trade *flows* or their changes to say something about implications for labor markets in sending and receiving countries

5. Rigorous way to think about economic consequences of trade deficits

6. Ability to distinguish trade from technology in driving empirical patterns in data

Distinct from H-O models that are based on differences in factor endowments with common technologies, the canonical conceptual model of international trade starts with Ricardo. Ricardo articulated the principle of comparative advantage: countries specialize in the activities in which they are relatively more productive. This productivity differences stem from differences in technology or skills (labor productivity). The Ricardian idea seems to many much more intuitively appealing than the H-O model, yet this idea had failed to have much traction in contemporary economic thinking until recently. Why? The *2012 Journal of Economic Perspectives* paper by Eaton and Kortum, "Putting Ricardo to Work," answers this question eloquently. A one word summary of their answer: tractability.

In its basic form, the Ricardian model produces knife-edge predictions–a series of special cases rather than general results. This has made the entire apparatus unattractive. The EK paper explains how economists have dealt with this problem historically and then lays out the solution that they and others have recently developed (their 2002 *Econometrica* paper is the classic reference). The solution, as we'll see, is not so much an "aha–that solves it!" as a workaround—a way to make the model tractable by choosing functional forms carefully and applying probability theory to get a continuum of results rather than a discrete number of special cases. Whether you find this approach appealing or merely a "trick" is somewhat a matter of personal taste. Inarguably, their innovations have cleared a logjam in applying the Ricardian model in both the theoretical and empirical realm, leading to a flowering of applied trade research that is arguably more relevant and appealing than the brand of trade economics that preceded it.

This section of the lecture note lays out the apparatus. You can get the same information from Eaton-Kortum's lovely paper at nearly the same level of detail. (I partly summarize the insights here for my own benefit. Note: if you are unsure whether you fully understand a topic, write a set of lecture notes on it. When you're done, you'll either understand the topic thoroughly or you'll be thoroughly clear that you don't get it.)

2.1 The basic Ricardian two-by-two model

Ricardo imagined two countries making two goods each. Let's take the case of Brazil and Costa Rica trading sugar and coffee. Assume that their labor requirements to make 100 kilos of each are:

	Coffee	Sugar
Brazil	100	75
Costa Rica	120	150

In this example, Brazil has an absolute advantage in both activities–that is, its unit labor requirements are lower in both sectors. However, it has a comparative advantage in Sugar since its unit labor requirements in Sugar relative to Coffee are 0.75 versus 1.25 in Costa Rica. Assume that the world relative price of coffee and sugar is 1

$$\frac{P_s^B}{P_c^B} = \frac{75}{100} < \frac{P_s^{CR}}{P_c^{CR}} = \frac{150}{120}.$$

Clearly, Brazil will export sugar and Costa Rica will export Coffee. Thus, Brazil can get 100 kilos of Coffee with only 75 units of labor instead of 100 if produced domestically, and Costa

Rica can get 100 kilos of sugar for only 120 units of labor rather than 150 units if produced domestically.

Even this simple example is obviously incomplete. Three types of outcomes are possible in this setting: (i) Brazil only produces sugar and Costa Rica only produces coffee; (ii) Brazil makes only sugar and Costa Rica produces both goods; (iii) Costa Rica makes only coffee and Brazil produces both goods. With incomplete specialization, relative prices of goods also have to meet market clearing conditions within a country–that is, if Brazil makes both goods, then the marginal product of Brazilian labor must be equated between coffee and sugar production. Once those prices are pinned down, we have to check whether consumer demands are consistent with market clearing. If not, we've got to check alternative cases. Even at this level of simplicity, the model is clumsy. One might speculate that it's not going to get prettier when we add more goods and more countries.

2.2 The chain of comparative advantage

Rather than assuming that the price of sugar and coffee are equated, let's normalize Brazil's wage to 1 and then determine the wage in Costa Rica, w, that is consistent with equilibrium. Assuming that trade occurs, it must be the case that the world price of coffee and sugar will be equated in both countries in purchasing power terms. Let's write these as $p(c)$ and $p(s)$. Competition assures that costs are minimized, so it must be the case that

$$p\left(c\right) = \min\left\{120w, 100\right\}, \; p\left(s\right) = \min\left\{150w, 75\right\},$$

where 100 and 75 correspond to the export prices of Brazilian coffee and sugar respectively. For a country to be able to trade, it must also be a producer. Thus, each country must be the minimum cost producer of at least one good. Applying this logic to the numbers above, it's evident that w must be at least 0.5 (its sugar price equals Brazil's) but no greater than 0.80 (its coffee price equals Brazils). Since Costa Rica has an absolute disadvantage in both goods, it must also be the case that $w < 1$, but this is automatically satisfied by the prior inequalities.

Let's write unit labor requirements as $a_B\left(c\right), a_B\left(s\right)$ and $a_{CR}\left(c\right), a_{CR}\left(s\right)$ for Brazil and Costa Rica respectively. Since Brazil has a comparative advantage in sugar, we can write

$$\frac{a_B\left(c\right)}{a_B\left(s\right)} > \frac{a_{CR}\left(c\right)}{a_{CR}\left(s\right)}.$$

One can readily extend this logic to three-plus goods to form a so-called "chain of comparative advantage"—that is, a series of inequalities that express the comparative advantage in one

5

country relative to the other. An equilibrium is the w that breaks the chain such that one set of goods is produced in Brazil, another in Costa Rica. At most one good can be produced in common; this occurs if at wage w, Brazil and Costa Rica have identical costs for producing the marginal good j. Note that this set of inequalities is not by itself sufficient to pin down the equilibrium. One needs further assumptions on demands and endowments to close the model.

Figure 1 in the EK *JEP* paper provides a nice illustration of wage determination (w^*) in the reference country (here Costa Rica, in their paper, England) as a function of labor supply.

Figure 1
Wage Determination in the Many Good Model

English share of world labor
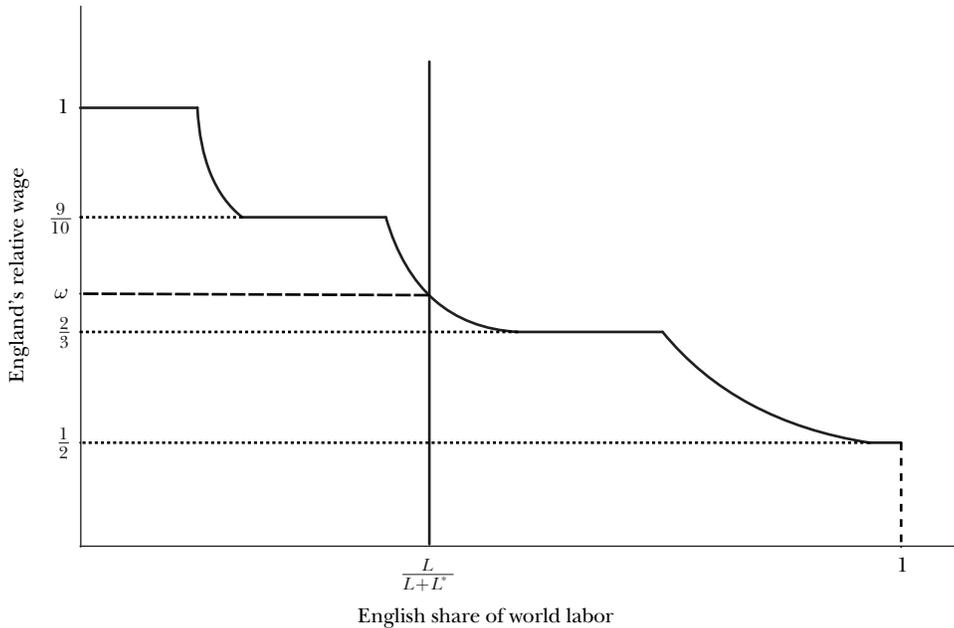
Note: The solid downward-sloping line is the relative demand curve for English labor, and the solid vertical line is the relative supply curve for English labor.

This figure makes a crucial point, which is that the demand curve with a finite number of goods will be decidedly non-smooth. In regions of the demand curve where the cost of the marginal good j is the same in the two comparison countries, demand for labor is perfectly elastic in each country since j can be produced in either country at the same cost. The conceptual experiment in Figure 1 of their paper corresponds to a case where England's (Costa Rica's) share of world labor supply expands such that it begins to take over production of additional goods. In the real where the two countries (England/Portugal, Costa Rica/Brazil) are producing the same good, demand for labor is perfectly elastic (no wage effects). When Costa Rica takes over production of the marginal good entirely–so the

two countries produce no goods in common, we are in a realm where demand is again elastic; the more of the original j good that Costa Rica produces, the more its price falls. If labor supply expands further, Costa Rica will eventually become competitive in the next good, j'. We are then in another flat spot where Brazil and Costa Rica are producing the (new) marginal good and the demand that each country faces is again perfectly elastic at the market price. It's easy to see why this demand structure will create problems for simple comparative static exercises.
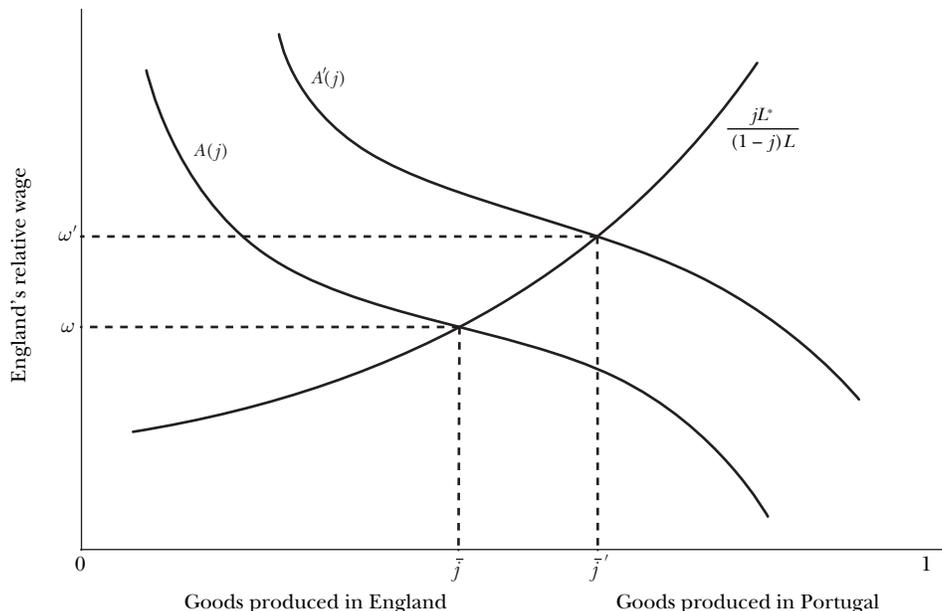
2.3   A continuum of goods

The problem with this simple model above is that demand for labor in each country is decidedly non-smooth as we transition between regions in which there is a marginal good j and no marginal good. A famous *AER* paper by Dornbusch, Fischer and Samuelson in 1977 (DFS) developed a nice workaround for this clunky structure. They assumed a continuum of goods j arrayed on the unit interval $j \in [0, 1]$ where the ratio $A(j) = a_B(j) / a_{CR}(j)$ is non-increasing in j and, more restrictively, $A(j)$ is smooth and strictly decreasing. Thus, Brazil's comparative advantage is rising in the index j. These assumptions guarantee that the chain of comparative advantage has no flat spots; there's always a marginal good that is equally costly to produce in both countries. In particular, continuing to use the normalization that the wage in Brazil is equal to one and the wage in Costa Rica is equal to w, the marginal good $\bar{j}$ satisfies

$$a_B(\bar{j}) = w \cdot a_{CR}(\bar{j}) \Rightarrow a_B(\bar{j}) / a_{CR}(\bar{j}) = w \Rightarrow A(\bar{j}) = w.$$

Thus, Costa Rica produces goods $j \leq \bar{j}$ and Brazil produces goods $j \geq \bar{j}$. A fall in w increases the range of goods that Costa Rica produces and contracts the range of goods that Brazil produces. Notice by the way that the Acemoglu-Autor (2010 Handbook chapter uses the same modeling tool. If you were to label the goods in DFS as tasks and the countries as skill groups, you'd have the foundation of the Acemoglu-Autor model (though what AA do with the model is distinct from DFS' agenda). Figure 2 of EK 2012 provides a nice illustration of the operation of the DFS model.

2.4   Trade costs

A second innovation of DFS is to introduce trade costs. These are missing from standard H-O models but arguably of first order importance in the real world. How do we know this is true? First, most countries consume a disproportionate share of their own output, which is consistent with either significant trade costs or strong home-bias in tastes. Second, distant countries trade less with one another and remote countries trade less with everyone. These facts suggest that trade costs have a large impact on equilibrium outcomes.

Trade costs also immediately introduce a greater semblance of reality into the frictionless H-O setting. If trade between two countries is costly, then competition is consistent with considerable heterogeneity in pricing. The same good can have different prices in different markets, and moreover, the low-cost producer of a good for one country may not be the low-cost producer for another country (if trade costs between these two different country pairs differ). Thus, trade costs are not just a nuisance, they're a crucial toe hold on reality.

DFS model trade costs as an "iceberg" phenomenon: part of the cargo decays (melts) in transit, and in general, the amount of decay will be proportional to transit time or distance.

Specifically, delivering one unit of a good from country i to k will require shipping $d_{ik} > 1$ units of the good. Note that d_{ik} will differ among country pairs. In general, it is assumed that d will not differ among goods *within* a country pair, but this can be relaxed. It's also typical to invoke the "triangle inequality," $d_{ik} \times d_{km} \geq d_{im}$, which says that the cost of shipping a good from country i to k and then from k to m is weakly greater than the cost of shipping it from i to m. This is a no-arbitrage condition.

2.5 Adding more countries

So far so good. But now it's about to get messy again. With many goods *and* many countries, the chain of comparative advantage is not likely to retain its useful 'smoothness' unless we impose strong restrictions on the shape of comparative advantage. (That's precisely what Acemoglu-Autor do for skill groups L, M and H. Arguably, it's justifiable to do this for skill groups, where productivity rankings might plausibly be monotone. For countries, this is harder to rationalize).

A classic paper by Jones (1961) shows why this doesn't work. The following table from Jones gives unit labor requirements for three goods in three countries:

	U.S.	Britain	Europe
Corn	10	10	10
Linen	5	7	3
Wool	4	3	2

Each country will produce one good, so the question is which? Normalize the U.S. wage at 1 and let w_B and w_E equal the wage rate in Britain and Europe respectively.

Here are two possible assignments

	Option 1	Option 2
Corn	Britain	U.S.
Linen	U.S.	Europe
Wool	Europe	Britain

You can check all of the inequalities for any two countries and two goods to confirm that both assignments satisfy comparative advantage. That is, for any pair of goods that two countries are producing, we can confirm that they would not want to switch places (so, in Option 1, we compare US and Britain in Linen and Corn, US and Europe in Linen and Wool,

and Britain and Europe in Corn and Wool).

$$\frac{a_{US}(C)}{a_B(C)} > \frac{a_{US}(L)}{a_B(L)} \Leftrightarrow \frac{10}{10} > \frac{5}{7}$$

$$\frac{a_{US}(W)}{a_E(W)} > \frac{a_{US}(L)}{a_E(L)} \Leftrightarrow \frac{4}{2} > \frac{5}{3}$$

$$\frac{a_B(W)}{a_E(W)} > \frac{a_B(C)}{a_E(C)} \Leftrightarrow \frac{3}{2} > \frac{10}{10}$$

For Option 2, we compare U.S. and Europe in Corn and Linen, U.S. and Britain in Corn and Wool, and Europe and Britain in Linen versus Wool:

$$\frac{a_{US}(L)}{a_E(L)} > \frac{a_{US}(C)}{a_E(C)} \Leftrightarrow \frac{5}{3} > \frac{10}{10}$$

$$\frac{a_{US}(W)}{a_B(W)} > \frac{a_{US}(C)}{a_B(C)} \Leftrightarrow \frac{4}{3} > \frac{10}{10}$$

$$\frac{a_B(L)}{a_E(L)} > \frac{a_B(W)}{a_E(W)} \Leftrightarrow \frac{7}{3} > \frac{3}{2}$$

Notice that both assignments satisfy Ricardo's inequality.

Which assignment is superior? Footnote 10 of Eaton-Kortum shows that efficiency demands the assignment that minimizes the *product* of labor requirements. For Option 1, that's $a_B(C) \times a_U(L) \times a_E(W) = 10 \times 5 \times 2 = 100$. For Option 2, that's $a_B(W) \times a_U(C) \times a_E(L) = 3 \times 10 \times 3 = 90$. So, the second assignment is cost-minimizing.

Clearly, this is iterative approach is messy: 1) test all feasible assignments to determine which satisfy comparative advantage; 2) calculate labor requirements; 3) choose the cost-minimizing feasible assignment.

Here's where the Eaton-Kortum 2002 *Econometrica* paper comes in. Their paper takes the next logical step–akin to DFS 1977–in turning a messy discreet problem into a tractable, continuous problem.

3   A probabilistic approach

One solution to the many-goods, many-countries problem would be to assume a continuum of countries, as DFS did when assuming a continuum of goods. But this is a bit silly. While there are arguably an arbitrarily large number of goods that can be traded, but there is only a finite, countable set of countries (actually, 196 at last count). In place of a continuum, EK assume a probability distribution and use then use properties of distributions (and the law of large numbers) to get predictions. Here's how they do it:

Continue to assume a continuum of goods $j \in [0,1]$ as in DFS and add an integer number of countries $i = 1, 2, ..., I$. We'll continue to work with only Labor as an input for now. Write the unit labor requirements for good j in country i as $a_i(j)$.

Here's where the cleverness starts. Think of these $a's$ as random draws from particular probability distributions. Using continuous distributions makes the problem smooth again. In addition, we don't have to keep track of individual values of the many $a's$ as we work through the problem; we can simply use cumulative distribution functions to determine the probability that a country is the low cost producer of a specific good.

The choice of functional forms is quite important here. We need a distributional family where conditional distributions inherit the property of marginal distributions, which moves us towards the exponential family (and not any exponential will do). The family that EK use is the Fréchet distribution. Why this distribution? If inventions (productivity realizations) are drawn from a Pareto distribution and only the most efficient draws are retained, then the order statistics from these draws will be characterized by the type II extreme value (Fréchet) distribution.[1]

Specifically, their assumption is that

$$\Pr[a_i(j) < x] = 1 - e^{-(A_i x)^\theta}.$$

Here, A_i is country $i's$ absolute advantage in all goods: a higher A_i means country i will in expectation have *lower* draws of labor requirements in all goods j. The parameter $\theta > 1$ is inversely related to the variance of labor requirements. A larger θ means that a country's labor requirements are typically closer to its country-specific mean; a smaller θ means that these requirements have greater dispersion.

A convenient simplification in this model (relaxed in more recent variants) is to impose the assumption that θ is common across countries. If so, θ also describes the type of competition that producers will face. If θ is large, so there's very little variance in productivity draws, what will matter is a country's absolute advantage as well as its trade costs vis-a-vis other countries. Competition in this setting is close to perfect (up to trade costs). Small perturbations in a country's cost vis-a-vis its competitors could cause it to gain or lose the entire market for a specific good j. Conversely, when θ is small, productivity draws will be very disperse. Even countries with low absolute advantage may be the low-cost (cheapest) producer of certain goods. Moreover, since the dispersion of countries' draws for a given good will be substantial,

---

[1]Following their innovation, the Fréchet distribution has become the foundation for models of comparative advantage. Arnaud Costinot has remarked that in every trade model, there comes a point when the author "Fréchets it up"—meaning, invokes the Fréchet distribution to get the functional forms needed to close the model.

small perturbations in costs may have little effect on the demand that a country faces for its output of good j. That is, if a country has strong comparative advantage in producing some good, then a modest increase in costs might not have a first order effect on its world market share for that good. In that case, competition has a strong component of comparative advantage and a much weaker component of perfect competition.

3.1   The price distribution

As before, let w_i equal the labor cost in country i (no normalization is needed, since it's only relative $w's$ that matter) and continue to assume iceberg transport costs such that $d_{ii'} > 1$ and $d_{ii} = 1$. Thus, the cost of producing good j in country i and delivering it country n is

$$c_{ni}(j) = a_i(j)w_i d_{ni}.$$

The price that country n pays for j is of course the *minimum* of all prices available to it. That is,

$$c_n(j) = \min\{c_{ni}(j)\}.$$

Note that this does not mean that each country pays the minimum world price for each good. With non-zero trade costs that differ among country pairs, the lowest available price for a good j will vary across countries.

The distribution of the cost of good j produced in country i and offered in country n is given by

$$\Pr[c_{ni}(j) < c] = 1 - e^{-(cA_i/w_i d_{ni})^{\theta}}.$$

Note this is only the price that n faces for good j from country i. The cumulative distribution of prices for good j that country n faces across all supplier countries is

$$\Pr[p_n(j) < p] = 1 - \prod_i \Pr[c_{ni}(j) > p]$$

$$= 1 - e^{-(\bar{A}_n p)^{\theta}}$$

$$\text{where } \bar{A}_n = \left[\sum_{i=1}^{I}(A_i/w_i d_{ni})^{\theta}\right]^{\frac{1}{\theta}}.$$

This term $\bar{A}_n$ is a country specific purchase price parameter. In EK 2002, they explain its meaning as follows:

> The price parameter $\bar{A}_n$ is critical to what follows. It summarizes how (i) states of technology around the world, (ii) input costs around the world, and (iii)

geographic barriers govern prices in each country n. International trade enlarges each country's effective state of technology with technology available from other countries, discounted by input costs and geographic barriers. At one extreme, in a zero-gravity world with no geographic barriers ($d_{ni} = 1$ for all n and i), $\bar{A}_n$ is the same everywhere and the law of one price holds for each good. At the other extreme of autarky, with prohibitive geographic barriers ($d_{ni} \to \infty$ for $n \neq i$), $\bar{A}_n$ reduces to A_n/w_n, country n's own state of technology down-weighted by its input cost.

As we will see below, a higher value of $\bar{A}_n$ will correspond to a *lower* price index, meaning that a country's PPP is rising in its $\bar{A}_n$. All else equal, a higher value of θ (less cross-country productivity dispersion) means that $\bar{A}$ tends towards the highest value of A; because there is little dispersion in cross-country productivities, gains from trade come largely from variation in labor costs and geography. Conversely, a lower value of θ (more cross-country productivity dispersion) means that $\bar{A}$ tends towards the sum of the $A's$; b/c productivities across countries are more dispersed, a country gains from differences in productivity across countries, not simply from variation in their labor costs or distances.

Finally, we want to add preferences. Without any loss of generality, we'll consider the simplest case: preferences are symmetric Cobb-Douglas, with equal shares of income spent on all goods. In this case, the ideal price index is simply the geometric mean of the price distribution. So we can write

$$p_n = \frac{\gamma}{\bar{A}_n},$$

where γ is a constant.[2] *Lower* values of p_n mean higher purchasing power. Since $\bar{A}_n$ enters inversely, a higher value of $\bar{A}_n$ corresponds to *higher* purchasing power. Factors that raise $\bar{A}_n$ (lowering p_n) are:

- Higher own-productivity (A_i);

- Lower bilateral trade costs (d_{ni});

- Lower input costs (w_i).

- Lower θ, reflecting greater productivity dispersion across countries

Thus, the model nicely gives rise to differences in PPP as a function of the model's primitives.

---

[2] $\gamma = e^{-\epsilon/\theta}$ and ϵ is Euler's constant.

3.2 Equilibrium

There's a whole lot more one can do with this model. One thing we haven't done so far is close it. In a model where all income is labor income, this is straightforward. Let L_i be the labor endowment of country i. Then country $i's$ total income is

$$w_i L_i = \sum_{n=1}^{I} \pi_{ni} \left(w_n L_n + D_n \right) \tag{1}$$

where D_n is the trade deficit in county n (meaning that it spends in excess of its labor income) and π_{ni} is the share $n's$ consumption purchased from i. We derive an expression for π immediately below. Equation (1) is actually a system of I linear equations. It will generally have to be solved numerically.

3.3 Putting the model to work: Trade shares and 'gravity'

We can now derive a set of useful comparative statics for trade flows.

- The probability π_{ni} that country i is the lowest cost supplier of any specific good j to country n is

$$\pi_{ni} = \Pr\left[c_{ni}\left(j\right) = p_n\left(j\right) \right] = \left(\frac{A_i / w_i d_{ni}}{\bar{A}_n} \right)^{\theta}.$$

  We don't subscript this probability by j because the probability does not differ across goods. Notice that θ plays the role of an elasticity. Imports of goods from i to country n decrease with elasticity θ in response to a rise in w_i or d_{ni}. Notice that higher world productivity (net of trade costs) measured by $\bar{A}_n$ lowers the probability that country i is the low cost producer of j for country n. [Notice that $\partial \ln / \partial \ln d_{ni} = -\theta$, meaning that the distance elasticity of trade is increasing in the dispersion of productivity (lower θ).]

- With a continuum of goods, π_{ni} is also the share of all goods (on the continuum) consumed in n that are supplied by i.

- Since π_{ni} is the fraction of goods that country n buys from i, and since the expected price of goods does not vary by source (conditional on purchase), π_{ni} is also the fraction of n's expenditure spent on goods from i. Let X equal expenditure:

$$\pi_{ni} = \frac{X_{ni}}{X_n} = \left(\frac{A_i / w_i d_{ni}}{\bar{A}_n} \right)^{\theta} \tag{2}$$

- Some further manipulation provides a useful expression. Let Q_i equal total sales by exporter i. Country i's total sales to the all countries m is:

$$Q_i = \sum_{m=1}^{N} X_{mi} = \left(\frac{A_i}{w_i}\right)^{\theta} \sum_{m=1}^{N} \frac{d_{mi}^{-\theta} X_m}{\bar{A}_m^{\theta}}. \tag{3}$$

- We can use (3) to solve for $(A/w_i)^{\theta} = (X_{ni}/X_n) \times (\bar{A}_n/d_{ni})^{\theta}$. Substituting into (2) gives an expression for X_{ni}, which is the consumption of goods in n produced in i (AKA exports from i to n):

$$\begin{aligned}
X_{ni} &= \frac{X_n \cdot Q_i \left(\bar{A}_n d_{ni}\right)^{-\theta}}{\sum_{m=1}^{N} \left(\bar{A}_m d_{mi}\right)^{-\theta} X_m} \\
&= \frac{X_n \cdot Q_i \left(\frac{\gamma d_{ni}}{\rho_n}\right)^{-\theta}}{\sum_{m=1}^{N} \left(\frac{\gamma d_{mi}}{\rho_m}\right)^{-\theta} X_m} \\
&= \frac{X_n \cdot Q_i \left(\frac{d_{ni}}{\rho_n}\right)^{-\theta}}{\sum_{m=1}^{N} \left(\frac{d_{mi}}{\rho_m}\right)^{-\theta} X_m}
\end{aligned} \tag{4}$$

This central equation says that the consumption of goods consumed in n that are produced in i are:

- Increasing in i's total economic size, Q_i. This is consistent with the idea that countries buy more from larger countries.
- Declining in bilateral trade costs d_{ni}
- Rising in the importer's total purchases X_n
- In all cases, the geographic barrier d_{ni} between i and the importer is deflated by the importer's price level ρ_n. The lower are goods prices in the destination market, the more that geographic barriers between n and i reduce trade. That is, import costs matter more when the destination market is more competitive.
- The denominator of this expression, $\left(\frac{d_{mi}}{\rho_m}\right)^{-\theta} X_m$, is the size of each destination market m as perceived by i. Higher X_m means that i has a larger market into which to sell. Higher bilateral trade costs d_{mi} and a lower price level ρ_m in market m reduces i's sales into m. Thus, county $n's$ share of $i's$ imports is $n's$ share of $i's$ *effective world market*—that is, the size of the world market perceived by i, after accounting for transports costs to i and prices and market size in all other countries.

3.4 Why do they call it the 'gravity' model?

- Historically, the gravity model for trade among bilateral country pairs had been estimated something like this:

$$\ln\left(X_{ni}\right) = \beta_0 + \beta_1 \ln\left(M_n\right) + \beta_2 \ln\left(M_i\right) + \beta_3 \ln\left(d_{ni}\right) + e_{ni},$$

  where M_n, M_i are the economic 'masses' of country's n and i, d_{ni} is the trade cost (usually distance), and the expected signs of β_1 and β_2 are positive and the expected sign of β_3 is negative.

- While this estimating equation precedes the development of the E-K gravity model and its relative, equation (4) from E-K's model also suggests this relationship. If we substitute X_n and Q_i for M_n and M_i, we get:

$$\ln X_{ni} = \ln X_n + \ln Q_i - \theta \ln d_{ni} + \theta \ln p_n - \kappa,$$

  where $\kappa = \ln\left[\sum_{m=1}^{N}\left(\frac{d_{mi}}{p_m}\right)^{-\theta} X_m\right]$.

- Equation (4) also implies a relationship between trade flows and price differences:

$$\frac{X_{ni}/X_n}{X_{ii}/X_i} = \left(\frac{\rho_i d_{ni}}{\rho_n}\right)^{-\theta}.$$

  The left-hand side of this expression is country $i's$ share in $n's$ output relative to $i's$ consumption of its own output. This share is falling in the bilateral trade cost. The share is rising as country $n's$ PPP falls relative to country $i's$, that is, as ρ_i/ρ_n falls (intuitively, it's expensive to import from countries with high PPP). Higher values of θ magnify these effects; when competition is more intensive (price heterogeneity smaller), trade flows are more responsive to given trade costs and price differentials.

3.5 Why large countries trade less

- Quoting from EK's 2012 JEP paper: An implication of the framework is that, given technologies around the world, having a larger share of the world labor force may require a country to have a lower wage. In order to employ more labor with its given set of technologies, a country needs to sell more of the goods it currently produces (going down a riser) or to take over goods from other countries (reaching a lower step). The result holds even though technologies are constant returns to scale, because larger size

reduces the gains from trade.

- To me, this statement only makes so much intuitive sense. What is the evidence that larger countries have lower wages? What *does* seem intuitive here is that larger countries will be their *own* low cost providers for a larger set of goods—simply because if there is a minimum efficient scale for an industry (not something that exists in E-K), larger countries will be able to support a larger number of industries at minimum efficient scale.

- Perhaps a trivial point, but one reason that larger countries may tend to trade less as a share of GDP is because their physical size actually makes trade more costly—serving to amplify international trade costs. That is, trading between Nebraska and Japan may be significantly more expensive than trading between California and trade. The typical E-K model variant assumes that transportation costs are nil within a country ($d_{ii} = 1$). If these costs are substantial, it could be an additional reason why internal regions of a country find it cheaper to either produce for themselves or purchase from adjacent regions rather than sourcing abroad. (This argument says that border regions between countries will engage in more international trade than internal regions, which is almost surely true in general.)

3.6   Why trade deficits matter

- If all countries have identical, homothetic preferences, there is no reason for countries to consume goods in different proportions. But a robust feature of data on trade and production is that countries tend to buy disproportionately more goods from themselves.

- One could explain this fact in terms of the basic Ricardian model by assuming that, for example, the Portuguese like wine more than the English. But it would be coincidental if tastes always happened to align with comparative advantage. (An alternative explanation is endogenous preference formation, e.g. Icelanders learned to like Hákarl[3] because it was there. Speaking as a foreigner, this strikes me as the *only* possible explanation for this abomination of national cuisine.)

---

[3]From Wikipedia: Hákarl is a national dish of Iceland consisting of a Greenland shark (Somniosus microcephalus) or other sleeper shark which has been cured with a particular fermentation process and hung to dry for four to five months. Hákarl is an acquired taste; it has a very particular ammonia-rich smell and fishy taste. [In popular parlance, Hákarl is called fermented shark, though my former six old used to call it demented shark.]

- Likely, a more central explanation for why countries consume their own output disproportionately is that moving goods between countries is costly. Because of iceberg trade barriers, goods no longer cost the same in each location. What happens if we introduce a trade cost d to all goods? The trade cost creates a range of goods that are not traded as each country makes them more cheaply for itself. As long as d is not too big, there is still a range of goods that England makes for everyone and another range that Portugal makes for everyone.

- An important implication of the trade cost is that it introduces a relationship between any trade deficit that a country runs and its relative wage. A transfer from England to Portugal diverts spending away from the non-traded goods that England was producing for itself toward the production of those same goods in Portugal. As a consequence, the English wage falls, leading to an expansion of the range of goods that England exports and a contraction of the range that Portugal exports.

- *Note that in a world with no trade frictions, trade deficits would not have this effect. If England transferred money to Portugal in a world with no trade frictions, Portugal would spend this money on English goods in the same proportions that England would. Where trade frictions exist, however, money transferred to Portugal from England is less likely to be spent on English goods and more likely to be spent on Portuguese (and other nearby) goods.*

3.7   Gains from Trade

One nice application of the framework is to derive the gains from trade. Real income in country i is

$$\frac{w_i}{\rho_i} = \gamma^{-1} A_i \pi_{ii}^{-1/\theta}.$$

This expression says that a country's income is increasing in its absolute advantage A_i (you can substitute the term "number of ideas" for A_i if you like) and its declining in its home share in consumption (because it's gaining less from trade). Taking logs to simplify we get

$$\ln\left(w_i/\rho_i\right) = -\ln\gamma + \ln A_i - \frac{1}{\theta}\ln\pi_{ii}.$$

Note that if a country doesn't trade at all, $\pi_{ii} = 1$ and hence the final term in this expression is zero, thus the real wage is determined entirely by domestic productivity A_i. Let's say that $\theta \simeq 4$. That is, the elasticity of trade with respect to price of traded goods is equal to 4. So, if country $i's$ import share went from zero to 0.25, so π_{ii} declined from 1 to 0.75, then its

welfare would rise by $-\frac{1}{4}\ln 0.75 \simeq 0.072$, i.e., 7.2 log points or about 7.5 percent. This insight (with a lot more depth) is the basis of the highly influential 2012 *AER* paper by Arkilakos, Costinot and Rodriguez-Claire.

3.8 What is this model for?

Other than a lot algebra, what have we gotten out of this model? From my perspective, here are some of the key 'deliverables:'

1. A succinct, closed analysis of the operation of comparative advantage in a full GE setting. This is not a small accomplishment. Of course, it requires strong distributional assumptions. But it's likely that theorists will make progress in relaxing them.

2. A plausible rationalization of the four key trade facts set out above:

   - trade between countries diminishes with distance;

   - prices vary across locations, with greater price differences between countries that are further apart;

   - factor rewards *do not* appear to be equalized across countries;

   - countries' relative productivities clearly vary across industries.

3. Perhaps most empirically relevant: This model relates trade *flows* to their potential effects on destination markets. Here, trade flows (and their changes over time) are directly informative about changes in comparative advantage and transport costs. These changes both induce trade flows and affect the final demand for the output of each competing country; hence, they impact welfare in competing economies. In the standard HO model, however, it is only prices and not flows that are informative about the potential impact of trade on a given market.

In summary, this paper offers the rudiments of a conceptual toolkit for analyzing how trade affects trading economies. I would say that this toolkit was largely lacking in the HO model. (Though see Krugman 2000 for a justification of the use of factor contents, which was a labor tool applied to trade questions, which trade economists such as Leamer (2000) had previously condemned as irrelevant.) As one partial example of this toolkit, the Autor, Dorn, Hanson (*AER* 2013) paper uses the EK model to consider how an exogenous rise in productivity in one country (China) would affect the demand for goods produced in each competing local economy, which ADH operationalize as 'Commuting Zones.' (The derivation of the ADH approach is given below. It is a very simple extension

The EK (2002) paper is now a decade since publication, and there has been much subsequent progress. Some of the key papers are:

- Costinot, "An Elementary Theory of Comparative Advantage." *Econometrica* 2009.

- Arkolakis, Costinot and Rodriguez-Clare, "New Trade, Same Old Gains." *AER* 2012.

- Melitz, "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity." *Econometrica* 2003.

- Costinot and Vogel, "Matching and Inequality in the World Economy." *JPE* 2010.

- Burstein and Vogel, "Factor Prices and International Trade: A Unifying Perspective." Working Paper.

And these are just the theory papers. Important empirical papers in this literature include:

- Verhoogen, "Trade, Quality Upgrading and Wage Inequality in the Mexican Manufacturing Sector." *QJE* 2008.

- Topolova, "Factor Immobility and Regional Impacts of Trade Liberalization: Evidence on Poverty from India." *American Economic Journal: Applied Economics* 2010.

- Donaldson, "Railroads of the Raj: Estimating the Impact of Transportation Infrastructure." *AER*.

- Eaton, Kortum, Kramarz, "An Anatomy of International Trade: Evidence from French Firms." *Econometrica*.

While the literature is progressing rapidly, it is a literature that for the most part lacks a labor economist's sensibility: tight identification, a focus on labor market consequences (not just prices and goods), and attention to institutions and frictions. Indeed, the labor market is notably absent from much (but not all) of contemporary trade literature, including Eaton-Kortum 2002. This provides an opportunity for labor economists.

4 Using the Eaton-Kortum (2002) model to motivate a labor market approach

A simple but potentially useful application of EK to labor markets is found in Autor-Dorn-Hanson *AER* 2013. Their reduced-form EK application is motivated as follows:

- Let the demand for labor in industry j by region i be given by $L_{ij} = L^d(w_{ij}, Q_{ij})$, where w_{ij} is unit production costs and Q_{ij} is output.

- Note that you should think of an industry as containing a continuum of goods within it. Thus, all of our reasoning above about "shares" by countries applies to shares by industries within and across countries as well.

- For region i, sales to destination market n in industry j are a function of its technological capability (A_{ij}), unit production costs (w_{ij}), and bilateral trade costs (d_{nij}), as well as expenditure in destination market n for goods of industry j (X_{nj}).

- Technological capability, A_{ij}, is a parameter that determines the position of the distribution of firm productivities in an industry and region. Using the EK model, region i's sales in industry j to destination market n can be written as

$$ X_{nij} = \frac{A_{ij}(w_{ij}d_{nij})^{-\theta}}{\Phi_{nj}} X_{nj}, $$

where θ is a parameter describing the dispersion in productivity among firms and $\Phi_{nj} \equiv \sum_h A_{hj}(w_{hj}d_{nhj})^{-\theta}$ describes the "toughness" of competition in destination market n in industry j, reflecting production and trade costs in the locations that supply products to market n.

- Region i will capture a larger share of market n's purchases in industry j when it has high productivity, low production costs, and low trade costs relative to other suppliers. Define $\tilde{A}_{ij} \equiv A_{ij}w_{ij}^{-\theta}$ to be the cost-adjusted productivity of region i in industry j. Then, summing over destination markets for region i, its total output in industry j is

$$ Q_{ij} = \tilde{A}_{ij} \sum_n \frac{X_{nj}d_{nij}^{-\theta}}{\Phi_{nj}}. \tag{5} $$

- China will be among the countries with which each U.S. region competes in serving destination markets. When China's productivity expands or its foreign trade costs fall, it increases the value of Φ_{nj} in each destination market, diverting product demand away from U.S. regions that also serve these markets.

- To show this formally, consider the change in Q_{ij} that would result were China to experience exogenous productivity growth (i.e., an increase in T_{cj}, where c indexes China) or a reduction in trade costs, do, say, to China's accession to the WTO. The

direct effect of changes in China's productivity and trade costs on Q_{ij} is

$$\hat{Q}_{ij} = -\sum_n \frac{X_{nij}}{Q_{ij}} \frac{X_{ncj}}{X_{nj}} (\hat{A}_{cj} - \theta \hat{d}_{ncj}) \tag{6}$$

where $\hat{x} \equiv d\ln x$, X_{nij}/Q_{ij} is the share of exports to destination market n in region i's output in industry j, and X_{ncj}/X_{nj} is the share of imports from China in spending by destination market n in industry j.

- Equation (6) implies that the fall in region i's output in industry j is higher is cost-adjusted productivity growth in China $\hat{A}_{cj}$ and the larger is the reduction in trade costs facing China $\hat{d}_{ncj}$, where the impact of these shocks is larger the more dependent region i is on market n and the more important China is as a source of supply to market n. In applying equation 6, we're focusing on competition that CZs face from China in the U.S. market, thus limiting the summation above to $n = u$, that is, to outputs produced and consumed in the United States.

- In general equilibrium, changes in China's productivity and trade costs may also cause wages and other factor prices to change in the countries with which China competes. These changes in factor prices, in turn, may cause changes in aggregate spending by countries, as the effects of shocks to China reverberate through the global economy. Equation (6) thus shows only the *direct* effect of shocks to Chinese productivity and trade costs on the demand for output in region i, ignoring the indirect effects of these changes on factor prices and spending in region i and in other regions and countries.

The useful feature of this setup is that it provides a link between observed changes in *quantities* of goods imported from China and changes in the demand for the output of a local economy (e.g., a Commuting Zone). Thus, unlike the H-O models we considered in the earlier lectures, this model provides an empirical toehold for quantities of goods traded to local demand for labor. This linkage is almost non-existent in the H-O framework.

5 A Case Study in Trade Adjustment: The Opening of Hong Kong to Trade with China [Hsieh and Woo, 2005]

China opened its market to foreign investors in 1980. Hsieh and Woo assert (and I don't know the historical basis for this claim—could be true, but there's probably another example

somewhere) that this resulted in one of the largest cases of increased 'outsourcing' in world history. I put the term outsourcing in ellipses because I believe that the policy experiment and analytic lens used by this paper are more closely akin to traditional trade theory rather than anything specific to outsourcing. As Table 1 makes clear, there was a remarkable change in the composition of employment in Hong Kong between 1981 and 1991, with the manufacturing share of employment falling from 39.3 to 19.0 percent in one decade.

How should be expect the opening to trade with China to affect, in particular: Employment of skilled versus unskilled labor within services and manufacturing in Hong Kong? Size of the service and manufacturing sectors in Hong Kong? The 'return to education' in Hong Kong? It is probably most useful to think of this as a case where a small, skill-intensive closed economy (Hong Kong) opens bidirectional to trade with a large, unskill-intensive closed economy (China). Thus, the equilibrium price ratio in Hong Kong is likely to shift heavily towards the Chinese autarkic price ratio.

We want to consider four types of evidence:

1. The change in 'skill demand' due to the change in the relative size of Services versus Manufacturing

2. The change in skill usage within the Manufacturing sector

3. The correlation between 'outsourcing' within Manufacturing industries and changes in skill usage

4. The change in the return to education. [Q: Why should this one come last?]

5.1   The change in 'skill demand' due to the change in the relative size of Services versus Manufacturing

We can write the change in the employment share of 'skilled' workers using the following decomposition:

$$
\begin{aligned}
D_{t1} &= D_{t1}^S \cdot E_{t1}^S + D_{t1}^M \cdot E_{t1}^M, \\
D_{t0} &= D_{t0}^S \cdot E_{t0}^S + D_{t0}^M \cdot E_{t0}^M. \\
\Delta D &= D_{t1} - D_{t0}, \\
\Delta E &= E_{t1} - E_{t0}.
\end{aligned}
$$

$$\begin{aligned}
\Delta D &= D_{t1}^S E_{t1}^S + D_{t1}^M E_{t1}^M - D_{t0}^S E_{t0}^S + D_{t0}^M E_{t0}^M \\
&= \left(E_{t1}^S - E_{t0}^S\right) \cdot \left(D_{t0}^S - D_{t0}^M\right) + \left(D_{t1}^M - D_{t0}^M\right) E_{to}^M + \left(D_{t1}^S - D_{t0}^S\right) E_{to}^S.
\end{aligned}$$

The three terms in this equation are:

1. The change in skilled employment due to growth of services relative to manufacturing

2. The change in skilled employment within manufacturing industries

3. The change in skilled employment in services industries

The first term is the canonical channel by which trade impacts skill demands: Expansion of export sectors, contraction of import sectors.

A second major channel is outsourcing of intermediate inputs. Within every industry, but esp. manufacturing, there are likely to be a set of 'tasks' that can be done abroad more cheaply. When these tasks are outsourced, the 'left over' work may be more or less skill intensive. This impact is capture by a decomposition performed within manufacturing:

$$\Delta D^M = \left(D_{t1}^M - D_{t0}^M\right) = \sum_j \left(\Delta E_j^M \cdot D_j^M\right) + \sum_j \left(\Delta D_j^M \cdot E_j^M\right).$$

The first term measures the change in employment in manufacturing industries, weighting by initial skill intensity. The second term measures the changes in skill intensity within industries, weighting by initial employment share.

Table 3 shows the results of this exercise. Within manufacturing changes in skill input are approximately twice as large a contributor to changes in skill input (measured by wagebill share) as is reallocation of employment to services (despite the fact that services are initially more skill-intensive). This is quite striking given that manufacturing employment (as a share of total employment) falls by 50% over this period.

This evidence has significant limitations. These patterns could be consistent, for example, with any factor augmenting skill-biased demand shift in a closed economy. It would be helpful to have some *more direct* evidence. The authors turn to within-industry wage-bill share regressions, where the explanatory variables are, most importantly, a measure of outsourcing. Also included are measures of capital/output (for cap-skill comp), log output (for scale effects), pre-existing trends, and skill upgrading in other countries within the same industry (a proxy for 'global SBTC'). The estimating equation is:

$$\Delta D_{t,j} = \beta_1 \Delta Out_{t,j} + \beta_2 \Delta \ln\left(K_{t,j}/Y_{t,j}\right) + \beta_3 \Delta \ln Y_{t,j} + \beta_4 Time_t + e_{t,j},$$

where outsourcing is defined as either $d\ln[\text{imports} / (\text{imports} + \text{shipments})]$ or $d\ln[\text{imported intermediate inputs} / \text{imported intermediate inputs} + \text{shipments}]$. These models are run using OLS and also instrumented with start of period variables proxying labor-intensity or high-skill employment. Following trade opening, outsourcing growth should be relatively greater in industries were initially more labor intensive and relatively smaller in industries that were initially intensive in high-skill labor.

OLS and IV estimates both suggest that industries that outsourced more had a greater growth in skilled wage-bill shares. That IV estimates are larger than OLS estimates is somewhat surprising. Simultaneity bias would generally be expected to lead to OLS estimates that are smaller than IV estimates.

Figure 1 shows a striking rise in the 'return to education' in Hong Kong after 1981. Should we be surprised that the rise is not larger in manufacturing than in the aggregate economy?

Overall, this paper provides a number of pieces of indirect evidence suggesting that opening of trade to China was responsible for the dramatic changes in skill mix (ratio of skilled to unskilled labor in production within industries) and skill prices in the Hong Kong economy from 1981 to 1996. No piece of this evidence is decisive. But the array of facts is very much suggestive of the GE effect of a type of trade-opening explanation that we normally only study in blackboard exercises rather than in extant economies.

6   Grossman/Rossi-Hansberg, 2008 [for self-study]

The 2008 *AER* paper by Grossman and Rossi-Hansberg offers a model that is somewhat more closely attuned to the phenomenon of outsourcing as many lay people and some economists perceive it. In the model, workers of two types, H and L, each perform a continuum of 'tasks,' some of which are more suitable for outsourcing than others. Their model explores what happens when the cost of outsourcing tasks declines, leading to an increase in the extent of task outsourcing among one or both skill groups, though *not* a complete elimination of tasks performed by either skill group. This model should not be viewed as the last word on its topic (closer to the first). There are some key assumptions—not all of them transparent—that make this model tick. One feature that makes it potentially controversial is that embeds a Ricardian (comparative advantage) framework inside of a standard Hecksher-Ohlin framework. This makes the model a bit complicated and fairly non-standard. Whether this approach will prove to be a useful conceptual advance is not yet certain.

Firms in the Home country produce two goods, X and Y, where X is relatively skill-intensive. There are two factors of production modeled (though others can be viewed as operating in the background), H labor, which performs H tasks, and L labor which performs

L tasks. Production of each good requires a continuum of L and a continuum of H tasks. The measure of tasks of each type used in production of each good is normalized to 1. Each task within a continuum uses the same amount of the relevant type of labor if performed at home. That is, if L-tasks i and i' are undertaken at home in the course of producing good j, then firms use the same amount of domestic low-skilled labor to perform task i as they do to perform task i'. (This normalization is probably harmless since tasks could be made 'wider' or 'narrower' to guarantee that this condition holds.) If industries differ in their factor intensities, then it will be the case that they differ in the amount of L or H labor used to perform all tasks in the L or H continuum.

A key implicit assumption is that there is *no substitution* between tasks. Started more strongly, all tasks in a continuum are *perfect complements* since all must be performed for production to take place—so, each task must be performed at fixed intensity to produce a unit of output. This observation immediately implies that forces that reduce the cost of performing some subset of tasks within a continuum will, all else equal, raise demand for the other tasks in that continuum.

Specifically, in industry j, a firm needs a_{fj} units of domestic factor f to perform a typical f task once. Since the measure of f tasks is normalized to one, this means that a_{fj} is the total amount of domestic factor f needed to produce a unit of good j in the absence of offshoring. The fact that X is relatively skill intensive implies that:

$$\frac{a_{Hx}}{a_{Lx}} > \frac{a_{Hy}}{a_{Ly}}.$$

Offshoring takes a very simple form. Firms can undertake tasks at home or abroad. It is assumed (reasonably) that tasks that are more 'routine' or 'rules-based' are easier (less costly) to offshore. GRH assume initially that only L tasks may be offshored, though the entire analysis carries through in parallel for H tasks. Order the L tasks in an industry by $i \in [0, 1]$ so that the cost of offshoring is non-decreasing in i. One way to model this is that the unit labor requirements for tasks performed abroad are greater than or equal to unit labor requirements for tasks performed at home. In particular, if task i requires a_{Lj} units of labor if performed domestically, it requires $a_{Lj}\beta t_j(i)$ units of foreign labor where β is a shift parameter and $\beta t_j(i) \geq 1$ for all i and j and $t'(i) > 0$ (making this inequality strict simplifies things considerably).

A key modeling choice is to ask which industry, X or Y, finds offshoring less costly? GRH start from the logical baseline that the costs of offshoring are common for a given factor regardless of where it is employed. Thus both X and Y face identical outsourcing costs: $t_X(i) = t_Y(i) = t(i)$. Substantively, this is equivalent to assuming that the unskilled

tasks in X are no more skill-intensive than the unskilled tasks in Y. This may not be a good assumption, but it is probably a good place to start. When you draw the Lerner diagrams implied by this paper (the paper does not provide diagrams), you will see how the combination of perfect complementarity between tasks performed by a skill group *within* an industry and equip-proportionate declines in the cost of offshoring tasks by a skill group *between* industries leads to a familiar theoretical case, though arrived at by different means.

It is further assumed that a_{Lj} and a_{Hj}, the task 'intensities,' are endogenously chosen to minimize firms' costs of production given the constraint that the chosen combination of intensities yields a unit of output.

Let w and w^* be the home and foreign wage, respectively, of L workers. Suppose that

$$w > \beta t\left(0\right) w^*,$$

so that it is profitable to outsource some L tasks.

Let I index the marginal task performed at home, so that:

$$w = \beta t\left(I\right) w^*.$$

If goods are produced competitively, price must equal input costs, so:

$$p_j \leq w a_{Lj}\left(\cdot\right)\left(1 - I\right) + w^* a_{Lj}\left(\cdot\right)\int_0^I \beta t\left(i\right) di + s a_{Hj}\left(\cdot\right) + ..., \text{ for } j = x, y.$$

Here, s is the wage of H labor and the notation $a_{Lj}\left(\cdot\right)$ and $a_{Hj}\left(\cdot\right)$ is meant to stress the dependence of the $a's$ on market conditions and technology so that these values are optimally chosen. This expression can be usefully rewritten as:

$$\begin{aligned} p_j &\leq w a_{Lj}\left(\cdot\right)\Omega\left(I\right) + s a_{Hj}\left(\cdot\right) + ..., \text{ for } j = x, y, \\ \text{where } \Omega\left(I\right) &= 1 - I + \frac{\int_0^I t\left(i\right) di}{t\left(I\right)}. \end{aligned}$$

This expression is important because it expresses total labor costs for L tasks in relation to the cost of performing all tasks domestically.

- In particular, $w a_{Lj}\left(1 - I\right)$ is the actual domestic labor cost

- The average foreign labor cost is $w^* a_{Lj}\beta \int_0^I t\left(i\right) di$.

- Substituting $w = \beta t(I) w^*$, the foreign labor cost is

$$\frac{wa_{Lj}\left(\int_0^I t(i)\,di\right)}{t(I)}.$$

- Thus, total labor cost is:

$$wa_{Lj}\left(1 - I + \frac{\int_0^I t(i)\,di}{t(I)}\right).$$

- Logically, $\Omega(I) < 1$, which can be seen from the fact that $t(i)/t(I) < 1$ for $i < I$ (thus the integral of $t(i)/t(I)$ over $[0,I]$ is less than I). Notice that β does not directly enter the expression for $\Omega(\cdot)$, but it enters implicitly since I depends upon β (generally $\partial I/\partial\beta < 0$).

- Market clearing in the market for L and H implies that

$$a_{Lx}X + a_{L_y}Y = \frac{L}{1-I},$$
$$a_{Hx}X + a_{H_y}Y = H.$$

Thus, an increase in outsourcing I also has the effect of increasing domestic labor supply of L by the factor $1/(1-I)$.

Assume that households have identical, homothetic preferences in *all* countries, and take the high skilled good X as the numeraire good, so $p_x = 1$, and $p = p_y/p_x = p_y$. If the home country is small, p may be taken as parametric, otherwise not.

6.1 Comparative statics

Taking Ω, p and I as exogenous for the moment, one can totally differentiate the equilibrium conditions to obtain an expression for the log change in the wage of low-skilled labor. In particular, let's assume that there are only two factors and two sectors and both sectors are active (so we are in the cone of diversification). In this case, the four equilibrium equations are:

$$p = wa_{Ly}(\cdot)\,\Omega(I) + sa_{Hy}(\cdot)$$
$$1 = wa_{Lx}(\cdot)\,\Omega(I) + sa_{Hx}(\cdot)$$
$$L = (1-I)a_{Lx}X + (1-I)a_{L_y}Y$$
$$H = a_{Hx}X + a_{H_y}Y.$$

The log change in w is equal to:

$$\hat{w} = -\hat{\Omega} + \mu_1 \hat{p} + \mu_2 \frac{dI}{1-I}, \tag{7}$$

where μ_1 are the terms multiplying the price change and μ_2 are the terms multiplying the labor supply term. This expression shows three separate channels by which outsourcing may impact the low-skilled wage.

6.2 Productivity effect

Consider the first term:

$$
\begin{aligned}
\Omega(I) &= 1 - I + \frac{\int_0^I t(i)\, di}{t(I)} \\
\frac{d\Omega(I)}{dI} &= -1 + \frac{\partial t(I)^{-1}}{\partial I} \times \int_0^I t(i)\, di + t(I)^{-1} \times \frac{\partial}{\partial I} \int_0^I t(i)\, di \\
&= -1 - \frac{t'(I)}{t(I)^2} \times \int_0^I t(i)\, di + 1 \\
&= -\frac{\int_0^I t(i)\, di}{t(I)^2} t'(I) < 0
\end{aligned}
$$

It must be the case that $\partial I/\partial \beta < 0$, that is a fall in the costs of outsourcing raises the extent of outsourcing: $\partial \Omega/\partial \beta > 0$. But if outsourcing rises, $\hat{\Omega} < 0$, then the first term of (7) implies that the low-skilled wage rises. Why? It's easiest to see by assuming that prices are exogenous, so the only moving part is Ω and hence $dp = 0$. Rewrite the two price as equations

$$
\begin{aligned}
p &= w a_{Ly}\,(\Omega w/s)\,\Omega(I) + s a_{Hy}\,(\Omega w/s). \\
1 &= w a_{Lx}\,(\Omega w/s)\,\Omega(I) + s a_{Hx}\,(\Omega w/s).
\end{aligned}
$$

Here the parenthetical $(\Omega w/s)$ terms on the input coefficients are intended to stress the dependence of the $a's$ on relative wages. If prices are parametric, then $dp = 0$. Also, we can

use the envelope theorem to approximate $\partial a\left(w/s\right)/\partial w \approx 0$:

$$
\begin{aligned}
dp &= 0 = a_{Ly}\left(\Omega w/s\right)\Omega dw + a_{Ly}\left(\Omega w/s\right)wd\Omega \\
0 &= a_{Lx}\left(\Omega w/s\right)\Omega dw + a_{Lx}\left(\Omega w/s\right)wd\Omega \\
dw\cdot\Omega\left[a_{Ly}-a_{Lx}\right]+d\Omega\cdot w\left[a_{Ly}-a_{Ly}\right] &= 0 \\
\frac{dw}{w} &= -\frac{d\Omega}{\Omega} \\
\partial\ln w &= -\partial\ln\Omega \\
\hat{w} &= -\hat{\Omega}.
\end{aligned}
$$

We can get a little more information by using the fact that

$$
\begin{aligned}
w &= w^{*}\beta t\left(I\right), \\
\hat{w} &= \hat{\beta}+\hat{t}\left(I\right), \\
-\hat{\Omega} &= \hat{\beta}+\hat{t}\left(I\right)=\frac{d\beta}{\beta}+\frac{t'\left(I\right)}{t\left(I\right)}.
\end{aligned}
$$

Rewriting:
$$
\frac{dw}{w}=\frac{dw^{*}\beta t\left(I\right)}{w^{*}\beta t\left(I\right)}=\frac{w^{*}t\left(I\right)d\beta+w^{*}\beta t'\left(I\right)dI}{w^{*}\beta t\left(I\right)}=\frac{d\beta}{\beta}+\frac{t'\left(I\right)}{t\left(I\right)}\frac{\partial I}{\partial\beta}
$$

Now, we substitute vigorously to get an expression for $t'\left(I\right)/t\left(I\right)$:

$$
\begin{aligned}
\frac{dw}{w} &= -\frac{d\Omega}{\Omega}=\frac{\frac{\int_{0}^{I}t(i)di}{t(I)^{2}}t'\left(I\right)\frac{dI}{d\beta}}{\Omega} \\
\frac{\frac{\int_{0}^{I}t(i)di}{t(I)^{2}}t'\left(I\right)\frac{\partial I}{\partial\beta}}{\Omega} &= \frac{w^{*}t\left(I\right)d\beta+w^{*}\beta t'\left(I\right)\frac{\partial I}{\partial\beta}}{w^{*}\beta t\left(I\right)} \\
\frac{\frac{\int_{0}^{I}t(i)di}{t(I)^{2}}t'\left(I\right)\frac{\partial I}{\partial\beta}}{\Omega} &= \frac{d\beta}{\beta}+\frac{t'\left(I\right)}{I}\frac{\partial I}{\partial\beta} \\
\frac{t'\left(I\right)}{I}\frac{\partial I}{\partial\beta}\left(\frac{\frac{\int_{0}^{I}t(i)di}{t(I)}-\Omega}{\Omega}\right) &= \frac{d\beta}{\beta} \\
\frac{t'\left(I\right)}{I}\frac{\partial I}{\partial\beta} &= \frac{d\beta}{\beta}\left(\frac{\Omega}{\frac{\int_{0}^{I}t(i)di}{t(I)}-\Omega}\right).
\end{aligned}
$$

Now substitute for dw/w:

$$\frac{dw}{w} = \frac{t'(I)}{I}\frac{\partial I}{\partial \beta} + \frac{d\beta}{\beta} \Rightarrow \frac{dw}{w} = \frac{d\beta}{\beta}\left(\frac{\Omega}{\frac{\int_0^I t(i)di}{t(I)} - \Omega} + 1\right)$$

$$\frac{dw}{w} = \frac{d\beta}{\beta}\left(\frac{\frac{\int_0^I t(i)di}{t(I)}}{\frac{\int_0^I t(i)di}{t(I)} - \left[1 - I + \frac{\int_0^I t(i)di}{t(I)}\right]}\right)$$

$$\frac{dw}{w} = -\frac{d\beta}{\beta}\left(\frac{\int_0^I t(i)di}{(1-I)t(I)}\right)$$

$$\hat{w} = -\hat{\beta}\left(\frac{\int_0^I t(i)di}{(1-I)t(I)}\right) = -\hat{\Omega}.$$

So, the bottom line here is

$$\hat{w} = -\hat{\beta}\left(\frac{\int_0^I t(i)di}{(1-I)t(I)}\right) = -\hat{\Omega}.$$

Since $\hat{\beta} < 0$ implies that outsourcing costs fall, we see that the wage effect is strictly positive.

This is easiest to see in a Lerner diagram. The reduction in the cost of performing L tasks on the interval 0 through I is a like a factor augmenting technical change that raises the productivity of L labor in both sectors. However, the extent of savings is larger in the Y sector, meaning that the cost of producing Y falls by more. Thus, both sectors become more H intensive, and the relative wage ratio must shift favorably towards L so that the relative cost of producing a bundle of X and Y is unchanged (otherwise, they cannot both be produced). L labor will be freed from both sectors, and this can be accommodate by an expansion of the L intensive sector.

6.3 Terms of trade and labor supply effects

If prices were parametric, this would be the end of the discussion. But if the home country is large, or if the fall in β is not unilateral (so multiple countries experienced a fall in outsourcing costs at once), then there would also be an output expansion effect, leading to an adverse terms of trade effect. The price of Y would fall as relative supply increased, and this would reduce the unskilled wage through the standard Stolper-Samuelson channel. This is the $\mu_1\hat{p}$ term in the above equation. These two terms (productivity gain, terms of trade) would be present in any standard two-by-two trade environment where a factor augmenting technical change differentially raised the productivity of the L factor.

The peculiar nature of outsourcing creates a third effect. In addition to the employment reallocation that results from the rise in w (due to increased L productivity), there is a mechanical effect coming from direct labor displacement (seen in $dI/(1-I)$. This is not present in a standard factor-augmentation case. It is as if outsourcing is a 'machine' that increases the productivity of a subset $1 - I$ workers while directly displacing the other I workers. Clearly, the term $\mu_2 dI/(1-I)$ must also be negative.

Thus, a reduction in β that raises I can have a net negative or positive effect on the L wage. This effect will be positive in the case where a small open economy experiences a unilateral reduction in its outsourcing costs. Obviously, such a case is highly stylized, and it is exactly this type of conceptual exercise that is critiqued by Krugman in 2000 (though Grossman and Rossi-Hansberg are sensitive to this criticism and acknowledge that this case is not realistic).

Meanwhile, for H labor, the log wage change equation can be written:

$$\hat{s} = -\mu_3 \hat{p} + \mu_4 \frac{dI}{1-I}.$$

There is no direct factor-augmentation effect here ($\hat{\Omega}$). The price effect will have opposite sign for H than L labor in the two-by-two case. The labor supply effect will generally be positive for the H wage due to standard q-complementarity effects.

6.4 More realistic cases

The paper discusses numerous extensions that expand from this narrow special case.

- So far, the paper assumes that the extent of outsourcing is identical in X and Y. It should be apparent that, holding p constant, if outsourcing possibilities increase differentially in Y, this is better for the wages of L and if they increase differentially in X, this is better for the wages of H. Draw the Lerner diagram and you will see why.

- It should also be clear that once we leave the small country case, an expansion in Outsourcing will have an adverse relative price effect on p, which (all else equal) will lower w. This works to the benefit of S and against w. It's possible, however, for both H and L to gain if the price effect does not swamp the productivity effect.

- The labor supply effect is implicitly included in $\hat{p}$ in most cases since domestic factor supplies do not impact relative wages within the cone of diversification except inasmuch as they change world relative output. Thus, the labor supply effect becomes relevant

when L is additionally used in a non-traded activity. In this case, the labor-supply effect of an increase in outsourcing works to the benefit of H and against L.

- All results hold in mirror image for a case in which offshoring of high skill tasks becomes feasible or increases. One interesting addendum here is that it is reasonable to assume (if we are thinking about industrial economies) that Home is relatively skill intensive relative to the rest of the world. Thus, even an uniform rise in outsourcing of H and L tasks differentially increases output of the X good (since this is the good in which Home is specialized). This means that potentially both H and L can gain from increased productivity, but the terms of trade effect will augment the benefits for L and reduce the benefits for H.

6.5   Conclusions from Grossman/Rossi-Hansberg

Whether or not you find this paper persuasive, you should certainly understand the conceptual mechanism. Among the paper's virtues is that the conceptual model is able to nest a large number of interesting cases in a relatively parsimonious form. The trick here is the paper's integration of Ricardian and Hecksher-Ohlin frameworks, which allows considerable flexibility, though at the cost of some transparency.

A central, under-examined, assumption of the paper is that all L tasks within an industry are complementary, and similarly for H tasks. At some level, this assumption seems right: senior programmers at IBM still need to architect the software that WiPro coders develop in Bangalore; physically present customer service representatives still need to sell mobile phones at the local Verizon store, even if the tech support is done from the Philippines. Thus, so long as *some* tasks need to be performed at Home, there is a fundamental complementarity between domestic and offshored tasks.

However, one may legitimately ask if the tasks 'left behind' tend to favor a particular 'skill group' within the distribution of H or L workers. One might loosely suspect that the *low-skill* tasks within the set of H tasks and the *high-skill* tasks within the set of L tasks are offshored (i.e., low-level programmers and moderately-skilled customer support positions are outsourced; high-level programmers and low-skill in-person sales agents are retained). This observation would be in the spirit of the ALM task framework, which suggests that it is 'routine' tasks that are increasingly subject to automation and offshoring (with routine tasks typically performed by the least educated among the highly educated and the most educated among the low educated). If so, it may not be innocuous to assume, as in GRH, that workers within a skill group H and L are homogenous. In particular, the L workers displaced by offshoring may be harmed if relegated to performing the remaining L tasks (if

these are particularly low-skill tasks), and the H workers displaced may be harmed if they are relegated to performing the remaining H tasks (if these are particularly high-skill tasks). Incorporating these observations would require a model with a larger number of skill types, or one in which there are only two skill types but that members of each type have heterogeneous abilities in various tasks. Such a model would be more complicated to write. But one should bear in mind Einstein's razor (a contrast to Occam's razor), "Everything should be made as simple as possible, but not simpler."[4]

An interesting exercise is to consider the effect of outsourcing in the Acemoglu-Autor Handbook model, where outsourcing is a technical change or price change that causes a subset of tasks previously assigned to a domestic skill group (L, M, or H) to be supplied by foreign labor. It might be useful to simplify down to two skills groups in the AA model and, for consistency with GRH, assume that the tasks that are outsourced are exclusively from the L subset of the task continuum (that is, they lie below I_l at its pre-outsourcing value). How would outsourcing of these tasks affect the relative wage of H versus L? How would it affect the real wage of H? How about the real wage of L (harder)? Why (and in what ways) do the implications of the AA model differ from those of GRH for this case? Which assumption(s) account for these differences? Which set of predictions strikes you as more plausible and why?

---

[4]The actual quotation from Einstein's 1933 Herbert Spencer Lecture is, "It can scarcely be denied that the supreme goal of all theory is to make the irreducible basic elements as simple and as few as possible without having to surrender the adequate representation of a single datum of experience."

MIT OpenCourseWare
http://ocw.mit.edu

14.662 Labor Economics II
Spring 2015

For information about citing these materials or our Terms of Use, visit: http://ocw.mit.edu/terms.