TRADE AND WAGES, RECONSIDERED

Paul Krugman

February 2008

This is a very preliminary draft for the spring meeting of the Brookings Panel on Economic Activity. Comments welcome.

There has been a great transformation in the nature of world trade over the past three decades. Prior to the late 70s developing countries overwhelmingly exported primary products rather than manufactured goods; one relic of that era is that we still sometimes refer to wealthy nations as "industrial countries," when the fact is that industry currently accounts for almost twice as high a share of GDP in China as it does in the United States. Since then, however, developing countries have increasingly become major exporters of manufactured goods, and latterly selected services as well.

From the beginning of this transformation it was apparent to international economists that the new pattern of trade might pose problems for low-wage workers in wealthy nations. Standard textbook analysis tells us that to the extent that trade is driven by international differences in factor abundance, the classic analysis of Stolper and Samuelson (1941) – which says that trade can have very strong effects on income distribution -- should apply. In particular, if trade with labor-abundant countries leads to a reduction in the relative price of labor-intensive goods, this should, other things equal, reduce the real wages of less-educated workers, both relative to other workers and in absolute terms. And in the 1980s, as the United States began to experience a marked rise in inequality, including a large rise in skill differentials, it was natural to think that growing imports of labor-intensive goods from low-wage countries might be a major culprit.

But is the effect of trade on wages quantitatively important? A number of studies conducted during the 1990s concluded that the effects of North-South trade on inequality were modest. Table 1 summarizes several well-known estimates, together with one crucial aspect of each: the date of the latest data incorporated in the estimate.

For a variety of reasons, possibly including the reduction in concerns about wages during the economic boom of the later 1990s, the focus of discussion in international economics then shifted away from the distributional effects of trade in manufactured goods with developing countries. When concerns about trade began to make headlines again, they tended to focus on the new and novel – in particular, the phenomenon of services outsourcing, which Alan Blinder (2006), in a much-quoted popular article, went so far as to call a second Industrial Revolution.

Until recently, however, surprisingly little attention was given to the increasingly out-of-date nature of the data behind the reassuring consensus that trade has only modest effects on income distribution. Yet the problem is obvious, and was in fact noted by Ben Bernanke (2007) last year: "Unfortunately, much of the available empirical research on the influence of trade on earnings inequality dates from the 1980s and 1990s and thus does not address later developments." And there have been a lot of later developments.

Figure 1 shows U.S. imports of manufactured goods as a percentage of GDP since 1989, divided between imports from developing countries and imports from advanced countries. [1] It turns out that developing-country imports have roughly doubled as a share of the economy since the studies that concluded that the effect of trade on income inequality was modest. This seems, at first glance, to suggest that we should scale up our estimates accordingly. Bivens (2007) has done just that with the simple model I offered in 1995, concluding that the distributional effects of trade are now much larger.

[1] Throughout this paper, manufactured goods are defined using the NAICS classification. "Advanced countries" are defined as the OECD less Korea, Mexico, and Turkey; developing countries are everyone else.

And there's another aspect to the change in trade: as we'll see, the developing countries that account for most of the expansion in trade since the early 1990s are substantially lower-wage, relative to advanced countries, than the developing countries that were the main focus of concern in the original literature. China, in particular, is estimated by the Bureau of Labor Statistics (2006) to have hourly compensation in manufacturing that is equal to only 3 percent of the U.S. level. Again, this shift to lower-wage sources of imports seems to suggest that the distributional effects of trade may well be considerably larger now than they were in the early 1990s.

But should we jump to the conclusion that the effects of trade on distribution weren't serious then, but that they are now? It turns out that there's a problem: although the "macro" picture suggests that the distributional effects of trade should have gotten substantially larger, detailed calculations of the factor content of trade – which played a key role in some earlier analyses – do not seem to support the conclusion that the effects of trade on income distribution have grown larger. This result, in turn, rests on what appears, in the data, to be a marked increase in the sophistication of the goods the United States imports from developing countries – in particular, a sharp increase in imports of computers and electronic products compared with traditional labor-intensive goods such as apparel.

Lawrence (2008), in a study that shares the same motivation as this paper, essentially concludes from the evidence on factor content and apparent rising sophistication that the rapid growth of imports from developing countries has not, in fact, been a source of rising inequality. But this conclusion is, in my view, too quick to dismiss what seems like an important paradox. On one side, the United States and other advanced countries have seen a surge in imports from countries that are substantially poorer and more labor-abundant than the third-world exporters that created so much anxiety a dozen years ago. On the other side, we seem to be importing goods that are more skill-intensive and less labor-intensive than before. As we'll see, the most important source of this paradox lies in the information technology sector: for the most part there is a clear tendency for developing countries to export labor-intensive products, but large third-world exports of computers and electronics stand out as a clear anomaly.

One possible resolution of this seeming paradox is that the data on which factor-content estimates are based suffer from severe aggregation problems – that developing countries are specializing in labor-intensive niches within otherwise skill-intensive sectors, especially in computers and electronics. I'll make that case later in the paper, while admitting that the evidence is fragmentary. If this is the correct interpretation, however, the effect of rapid trade growth on wage inequality may indeed have been significant.

The remainder of this paper is in four parts. The first part offers an overview of changing U.S. trade with developing countries, in a way that sets the stage for the later puzzle. The second part describes the theoretical basis for analyzing the distributional effects of trade, then shows how macro-level calculations and factor content analysis yield divergent conclusions. The third part turns to the case for aggregation problems and the implications of vertical specialization within industries. A final part considers the implications both for further research and for policy.

The changing pattern of trade

Figure 1 showed the dramatic rise in U.S. imports of manufactured goods from developing countries since 1989. One qualification that needs to be made right away is that to some extent this rise reflects the overall movement of the United States into massive trade deficit. The theoretical analysis later in this paper suggests that the average of imports and exports may be a better guide to likely distributional effects than imports alone. Figure 2 shows this number for U.S. trade in manufactured with developing and advanced countries; the rise in developing country trade is slightly less dramatic, but still impressive. Also note that 2006 marked a watershed: in that year, for the first time, the United States began doing more overall trade in manufactured goods with developing countries than with other advanced countries.

This rapid growth in U.S. trade with developing countries mainly took the form of increased trade with countries that were only minor players in the early 1990s. At the time of the original literature on trade and income distribution, North-South manufactured trade was still, to a large extent, trade involving the original four Asian "tigers": South Korea, Taiwan, Hong Kong, and Singapore. Since then, however, U.S. trade growth with developing countries has principally involved China, Mexico, and some smaller players. Figure 3 is an area chart of U.S. manufactured imports from developing countries, again as a percentage of GDP; it shows a modest relative decline for the original tigers and a large rise for Mexico and especially China.

This changing direction of North-South trade has one immediate implication: the aspect of this trade that initially attracted so much (often hostile) attention – the fact that we were now importing manufactured goods from countries with low wages by advanced-country standards – is much more extreme now that it was in the early 1990s. In 1990, according to BLS estimates, the four original tigers had average hourly compensation in manufacturing equal to 25% of U.S. levels. By 1995 that had risen to 39% of U.S. levels. But as of 2005 the BLS estimated that Mexico had hourly compensation only 11% of the U.S. level, and China only slightly more than 3%.

As a result, one trend that was often cited in the early 90s as a reason to discount fears about the effect of trade on wages – the fact that the average wage of U.S. trading partners was actually rising relative to the U.S. level – has gone into reverse. Table 2 shows the top 10 U.S. trading partners and the average hourly compensation of manufacturing workers in that group, weighted by the value of bilateral trade and expressed as a percentage of the U.S. level, since 1975. This measure did indeed rise from 1975 to 1990, reflecting rising relative wages both in advanced-country trading partners and in the original Asian tiger economies. Since 1990, however, the rapidly rising weight of China and, to a lesser extent, Mexico has driven the index down by approximately 20 percent.[2]

What accounts for the rapid growth of manufactured imports from these new players? China's economy, at least, has grown very rapidly, and one might imagine that the growth of China's exports is simply a reflection of its overall growth. Simple gravity models, in which the trade between any pair of countries reflects the product of their GDPs, adjusted for the distance between them, generally work quite well and have become a standard tool for interpreting the overall pattern of trade. And such a model would lead us to expect U.S. imports from China as a percentage of GDP to rise, other things equal, in proportion to Chinese GDP as a share of U.S. GDP.

In fact, however, U.S. imports from China have risen much more rapidly than the growth of the Chinese economy, on its own, would have led us to expect. Table 3 compares the growth in Chinese and Mexican GDP as a share of US GDP with imports from each country as a

[2] The most commonly used measure of the relative wages of U.S. trading partners, from the Bureau of Labor Statistics (2006), looks somewhat different from Table 2: it shows a more rapid rise between 1975 and 1990, from 62 to 80, and no change from 1990 to 2005. However, the BLS measure is fixed-weight: hourly compensation in each country is weighted by 2004 trade with the United States. As a result, the BLS index does not reflect the shift of U.S. manufactures trade to developing countries.

percentage of US GDP. Chinese GDP, at market exchange rates, has tripled relative to the United States – but US imports of manufactured goods from China have increased more than eightfold as a share of GDP. Mexico's GDP as a share of US GDP has risen about 40 percent, but manufactured exports have tripled relative to US GDP.

The obvious explanation of this "excess growth" in manufactured exports is that it reflects reduced barriers to trade, which have led to greater international specialization and hence greater trade. In the case of Mexico, it's natural to guess that NAFTA has played an important role, although much of the growth in Mexican exports may also reflect two other factors: the delayed effects of Mexico's dramatic unilateral liberalization of trade between 1985 and 1988, and the weak peso that followed the 1994-5 financial crisis.

In the case of China, there is no comparable break in policy. However, work by Hummels, Ishii and Yi (2003) suggests that even modest declines in trade costs can lead to large increases in the volume of trade by encouraging vertical specialization – the breakup of the production process into geographically separated stages. Thus rapid growth in Chinese exports might reflect declines in the cost of international communication and shipping.

One piece of evidence that may support the view that rapid growth in imports from developing countries reflects declining trade costs, both explicit and implicit, is the changing composition of these imports. A quick way to see the extent of this change in composition is to rely on a distinction introduced by Faberman (2004). In analyzing job loss and gain he distinguishes a group of "trade sensitive" industries (at the NAICS three-digit level) with very large import shares that also corresponds quite well to goods that we traditionally associated with third-world exports. Figure 4 shows the long-term trend in U.S. imports of manufactured goods from developing countries as a percentage of GDP, divided between “trade sensitive” and other goods. Even in 1989, it turns out, traditional third-world manufactured exports accounted for less than half of U.S. imports from developing countries. More to the point, however, the bulk of the growth in imports since then has come from non-traditional sectors.

What are these non-traditional goods? Figure 5 shows the change in imports from developing countries as a share of GDP by three-digit NAICS sector, from largest to smallest. The striking point is, of course, the extraordinary growth in imports of computers and electronics.

Modeling the effects of trade on income distribution

There have been two major waves of innovation in international trade theory over the roughly 30 years since developing-country exports of manufactured goods began to be a significant concern: the increasing returns/imperfect competition revolution of the 1980s and the more recent focus on intrafirm differences in productivity and propensity to export within industries. It is not clear, however, how to apply the insights of either set of ideas to the question of distributional effects of developing-country exports. As a result, most analysis of this issue continues to rely on the simple perfectly competitive factor-proportions model.

The first key insight from this model is the Stolper-Samuelson relationship between goods prices and factor prices. Consider a world in which there are two factors of production, skilled labor and unskilled labor, and two goods produced competitively under constant returns to scale, a skill-intensive good X and a labor-intensive good Y. Assume that workers move freely between firms and industries, so that all workers of each type receive the same wage. Finally, assume provisionally that an economy produces both goods. Then there is a one-to-one relationship between the relative prices of the two goods and the relative wages of the two types of labor. Letting a "hat" represent proportional rate of change,

$$\hat{P}_X - \hat{P}_Y = (\theta_{SX} - \theta_{SY})(\hat{w}_S - \hat{w}_U)$$

where θ_{SX}, θ_{SY} are the shares of skilled labor in the production cost of X and Y respectively.

Figure 6 completes the story. The left panel shows the relationship between relative goods prices and relative factor prices. The right panel shows the relationship between factor prices and the ratio of skilled to unskilled labor used in production. In each industry, a rise in the relative wage of skilled workers leads to a fall in the ratio of skilled to unskilled workers. This is one way to see the logic behind the Stolper-Samuelson result. As long as the country continues to produce both goods, a rise in the relative price of the skill-intensive good must lead to a rise in the relative wages of skilled workers. This implies a fall in the ratio of skilled to unskilled workers in *both* industries – and hence a fall in the marginal productivity of unskilled workers in terms of both goods. And that, in turn, means that the real wage of unskilled workers unambiguously falls.

It's worth noting one more point about this analysis: the Stolper-Samuelson process involves a complex reshuffling of resources between industries – sort of a swing-around-and-change-partners move. Consider what happens, according to this model, if there is a rise in the relative price of X. Production within each industry becomes less skill-intensive, yet overall employment of both factors remains unchanged because the industrial mix of production shifts toward skill-intensive industries. This is not a process one should expect to play out in full in the short run; the moral I would take from this is that Stolper-Samuelson should not be taken too seriously when interpreting data over short periods, say 5 years.

But the focus of this paper is on a somewhat longer period – the years since the early 1990s, whose data were the basis for the relatively benign estimates of the effect of trade on wages that still dominate discussion. Are the data since then consistent with a strong Stolper-Samuelson effect?

At first glance, the answer appears to be yes.

Consider first how prices have changed. The Bureau of Labor Statistics offers data on the prices of manufactured imports from developing and advanced countries. If we assume that developing countries export labor-intensive goods to the United States while advanced countries export skill-intensive goods (an assumption that we'll confirm with a major asterisk – the case of computers and electronics -- in the next section), the ratio of these prices should offer a measure of the relative price of labor-intensive goods. Figure 7 shows the log of this ratio, normalized so that 1995=0. It seems that there has indeed been a substantial decline in the relative price of labor-intensive goods since the mid-1990s.

Consider next changes in relative wages. Figure 8 shows two widely used indicators of wage differentials: the 90-50 ratio of hourly wages, and the college-noncollege ratio. Both are shown for men only, to abstract from changes in sex differentials; both are also expressed in logs, normalized so that 1995=0. Both measures have risen substantially since 1995.

It should be pointed out that Lawrence (2008) reaches a different conclusion, arguing that trends in relative wages are *not* consistent with a trade-driven story. This different in interpretation arises, I believe, from two factors. First, Lawrence uses earnings data aggregated across sexes, which does not show as strong a rise in inequality as the male-only data. Second, he focuses on the period since 2000 rather than the longer stretch since the mid-1990s.

I would argue that this short-term focus is problematic in two respects. First, on general principles it's not clear what one learns from very short-term movements in relative wages. As argued above, the implied adjustment from Stolper-Samuelson involves a complex reallocation of resources across industries, making it unsuitable for short-term analysis. Second, more specifically, the period since 1995 includes a major boom-bust cycle in high-technology industries. The technology bubble of the late 1990s probably elevated the education premium, while the subsequent bust caused that premium to deflate. As a result, inferences from the movement in inequality during the first few years after the tech bust should be taken with a grain of salt.

Perhaps the more general point is that Stolper-Samuelson is a ceteris paribus proposition, and cannot be refuted – or, to be sure, confirmed – based on the movement of relative wages alone.

That said, the combination of the price changes shown in Figure 9 and the wage changes shown in Figure 10 does look reasonably supportive of the proposition that rapid growth in North-South trade since the studies of the mid-90s has made the effects on inequality substantially larger. There is, however, a big problem with that conclusion: the result of using the methods I and others applied to the subject of trade and wages in the 90s to more recent data does not, at least on first appearances, fit the story.

There was a fairly heated dispute in the 1990s over the appropriate way to analyze the effects of North-South trade on wages. Some economists, notably Leamer (1994) argued that since the relationship shown in Figure 6 is between goods prices and factor prices, the only legitimate approach is to rely on price information, rather than on the volume of trade, which is endogenous. Others, myself included, argued that this represented a confusion between the question of how best to present models with the question of constructing the appropriate thought experiment for analysis – it makes sense to present Stolper-Samuelson as a goods-price-factor-price relationship, but in the real world prices are as endogenous as trade volumes. The appropriate method, I argued in Krugman (1995), was "but for" analysis: compare goods and factor prices with an estimate of what they *would have been* but for the opportunity to engage in manufactures trade with developing countries. And this but-for analysis inevitably leans strongly on calculations involving trade volumes.

Figure 9 illustrates the thought experiment from Krugman (1995). As in Figure 6, we assume that there are two goods, one skill-intensive, one labor-intensive. PPF represents the production possibilities of an aggregated OECD. If it were not possible to trade skill-intensive goods for labor-intensive imports from developing countries – that is, *but for* the possibility of North-South trade – equilibrium would be at the autarkic point A. In fact, however, this possibility exists; the opportunities for trade with newly industrializing economies are represented by the offer curve NIEO. As a result, equilibrium production is at Q while equilibrium consumption is at C, with the line PP representing relative prices in trade. The relative price of skill-intensive goods is higher, and that of labor-intensive goods lower, than would obtain in the absence of trade. Hence, the Stolper-Samuelson effect applies.

In the original analysis I created an extremely simple CGE model to calculate a back-of-the-envelope estimate of this but-for effect. This appears at first sight to be a significantly different approach from analyses such as that of Borjas, Freeman, and Katz (1997), which instead try to calculate the factor content of trade – the factors of production embodied in imports and exports. However, in later work (Krugman (1996), Krugman (2000)), it became apparent that the factor content approach, interpreted carefully, is fully consistent with an analysis based on trade flows and their effect on relative prices.

Figure 10 shows how this reconciliation can be carried out. Imagine holding goods prices constant while altering the economy's factor endowment, subtracting skilled labor while adding unskilled labor. This would have the effect of shifting the production possibility frontier inward at the lower right, but upward at the upper left, as illustrated by the shift from the production possibility frontier PPF_1 to PPF_2. Production would also shift, at constant goods prices, toward less output of the skill-intensive good and more output of the labor-intensive good. More specifically, suppose that at the initial factor prices the value of factors added is equal to the value of factors subtracted. Then production would shift northwest up the relative price line PP. If the change in factor endowments is sufficiently large, production will reach point C – that is, production will match consumption, so that trade is eliminated. And what is this change in factor endowments? It is precisely equal to the factor content of the initial volume trade, as measured using the factor content of each good's production per dollar of value in the advanced economy (*not* in the developing country) under the actual trading regime.

Now, imagine that after we have added the factor content of trade to the advanced economy, thereby eliminating the need for actual trade, we do a further thought experiment, in which we do two things: first, we eliminate the possibility of trade, and then we reverse the change in factor endowments. This would shift the production possibility frontier back to PPF_1, but because trade is no longer possible, consumption, production and relative prices would end up at the original autarky point A.

This may seem rather roundabout, but what it says is the following: the but-for thought experiment of eliminating North-South trade has the same effect on wages as another thought experiment in which we take a non-trading economy whose resources include the actual economy's factor endowment plus the factor content of the real economy's trade, and then eliminate that difference in factor endowments. In this sense, then, the factor content approach, carefully interpreted, is equivalent to the but-for trade analysis.

There are two advantages to thinking about the issue in terms of factor content. One is that it simplifies the interpretation of any structural model. In general, the results of any such model depend on all the parameters – factor shares in production, goods shares in spending, and all the relevant elasticities of substitution in both production and consumption. However, thinking in terms of factor content makes it clear that these parameters matter only insofar as they affect one derived number, the aggregate elasticity of substitution between skilled and unskilled labor. This simplifies sensitivity analysis, and in general helps clarify interpretation.[3]

The other advantage of thinking in terms of factor content is that it simplifies the task of empirical work – or at least that it seemed to do so in the past. Rather than having to calibrate a full model, the researcher can simply estimate the factor content of trade, which is informative in itself, and assess likely impacts by examining the implications of alternative aggregate elasticities of substitution.

All of this assumes that we can do a reasonably good job of measuring factor content. Before we get there, however, it's useful to extend the analysis to allow for an important feature of U.S. trade, especially recently: large trade deficits.

[3] It's important to note, however, that the relevant elasticity of substitution between factors is the one that would prevail in the absence of trade. So it is somewhat problematic to rely, as for example Borjas, Freeman, and Katz (1997) do, on estimates of this elasticity from time series that include a period of significant trade.

Figures 9 and 10 are real-trade theory diagrams, assuming, as must be the case under standard real-trade models, that trade is balanced. Clearly that's not a reasonable assumption for the United States today, which runs large trade deficits financed by capital inflows. (Circa 2005 a rough description of the U.S. economy was that we made a living selling each other houses, paying for them with money borrowed from China.) However, it's possible to use the factor-content approach under conditions of trade deficit provided that we make two further assumptions. First, we assume that the effects of capital inflows on demand are equivalent to a transfer payment made to domestic households. (That is, capital inflows are spent in the same way as domestically earned income.) Second, we assume that all domestic consumers have identical homothetic preferences – so that the composition of spending does not depend on who is receiving income.

Under these assumptions, the factor-content exercise can be represented by Figure 11. Here the economy's actual production and consumption are once again at Q and C, but this time the value of consumption at world prices PP is greater than that of production. The difference is the trade deficit, represented as a transfer of income to domestic consumers. Again, it's possible to construct a hypothetical economy that would produce the actual economy's consumption without the need for trade; this can be done by adding the actual factor content of trade to the original economy, which shifts the production possibility frontier from PPF_1 to PPF_2. The effect of trade on factor prices can then be inferred by subtracting the factor content of trade out again. Because of the assumption of homothetic preferences, the effect on relative factor prices depends on the extent to which the ratio of factors is altered in this exercise. In particular, even if a country runs so large a trade deficit that it is implicitly an importer of both skilled and unskilled labor, trade still raises the skill premium as long as constructing the hypothetical no-trade economy requires increasing the quantity of unskilled labor by more, in percentage terms, than the quantity of skilled labor.

And now we get to the fundamental empirical puzzle. In the early to mid 1990s, factor content exercises indicated a significant though modest move in the expected direction. The most recent estimates, however, suggest that the dramatic expansion of imports from low-wage countries since 1990 has not significantly enlarged the factor content of trade.

Table 4 shows estimates of "job displacement" by education level – another name for factor content -- from Mishel et. al (2007). The estimates were constructed using changes in the ratios of imports and exports to sales within each 4-digit NAICS industry to estimate changes in sales due to trade; these estimates were then run through the input-output tables to estimate total implied changes in output; finally, estimates of college and noncollege workers per unit of output from U.S. data were used to estimate factor content.

What the estimate shows is that rising trade deficits have made the United States a consistent importer of both highly educated and less educated labor – that is, the U.S. picture looks like Figure 11, with factor content arising from a trade deficit as well as comparative advantage, rather than Figure 10. Nonetheless, prior to 1989 the estimated effect of trade was a relative increase in the effective supply of less-educated labor.

Since then, however, the calculations of Mishel et. al. indicate little net effect of trade on relative effective factor supplies. The obvious explanation lies in the trends illustrated in Figures 4 and 5: although the traditional manufactured exports of developing countries to the United States are labor-intensive goods like apparel, the growth in developing-country exports has been concentrated in nontraditional sectors, especially computers and electronics.

As we'll see next, the apparent strong comparative advantage of developing countries in these industries seems anomalous – unless the exports of developing countries are concentrated in labor-intensive sub-sectors within the industries.

Within-industry specialization and the problem of interpretation

A useful overview of the seemingly anomalous nature of some developing-country exports can be obtained by using a technique suggested by Romalis (2004). Romalis provided impressive evidence of the continuing relevance of Heckscher-Ohlin trade theory based on an analysis of the sources of U.S. manufactured imports. He showed that the United States does tend, systematically, to import skill-intensive goods from advanced countries and labor-intensive goods from developing countries, although the relationship is far from perfect – an imperfection he ascribed to the interaction of product differentiation and transport costs, as modeled in Krugman (1980). An alternative interpretation, of course, is that the evidence is blurred by measurement error.

Figure 12 is a simple Romalis scatterplot, deliberately conducted at a relatively high level of aggregation. The data points are 3-digit NAICS industries. The horizontal axis shows skill intensity as proxied by the share of nonproduction workers in employment. The vertical axis shows the share of developing countries in U.S. imports within the industry. It is immediately apparent that most industries fall along a downward-sloping "main sequence" in which developing countries tend to export labor-intensive goods, with apparel and other traditional third-world exports at the upper left.

I have identified the industries that lie clearly off this main sequence. The industries at the lower left pose little puzzle: the paper and wood products industries aren't very skill-intensive, but U.S. imports within these industries are, for reasons of resource abundance and geography, dominated by Canada. (They have lumberjacks, and that's OK.)

But what are we to make of NAICS 334, Computer and Electronic Products? In U.S. data it ranks as the most skill-intensive of industries, yet it is also an industry in which more than three-quarters of imports come from developing countries, especially China.

It seems a foregone conclusion that aggregation is a serious problem here; why not use more disaggregated data? The answer is that within the limits of what is possible given current data, we can do little to resolve the issue.

First of all, factor content analyses are limited by the level of disaggregation of the input-output table, which is at 4-digit level. A 4-digit version of Figure 12 looks essentially the same: all the components of 334 remain in the upper right corner. And even finer levels of disaggregation are of relatively little help. To see why, look at Table 5, which shows the five 6-digit sectors with the largest U.S. value-added within Computer and Electronics; collectively these sectors account for 57 percent of the total.

These are not homogeneous sectors. They are, however, globalized industries – and it is easy to find qualitative information suggesting that there is a division of labor between skill-intensive operations and less skill-intensive operations within each industry.

This is obvious in the case of computers. First of all, there is a clear division between the types of computers produced in emerging Asia – primarily relatively low-end, standardized products – and those produced in advanced countries. Probably even more important, computer production involves many stages, which are commonly split between advanced and developing economies in a way clearly related to skill-intensity. This paper was written on a Lenovo notebook computer. Lenovo, which took over the ThinkPad line from IBM, is famously a Chinese firm whose headquarters and design operations are in North Carolina, and many of whose components are produced in advanced countries. These operations help make the computer industry look highly skill-intensive, if we rely on data from the U.S. Census of Manufactures; this is not a good representation of what the industry looks like in China.

The caricature of the computer industry is that Japan and the U.S. make the innards, then China adds the plastic shell. While it isn't actually that simple, Dean, Fung and Wang (2007) estimate that imported inputs account for 52 percent of the value of Chinese computer exports. Similarly, imported inputs account for 41 percent of electronic device value, 46 percent for electronic appliances, and 59 percent for communications equipment. And there is little question that in each case the imported inputs are much more skill-intensive than the Chinese component of the process.

Semiconductors might seem like a more homogeneous product. But even the semiconductor industry is marked by an international division of labor that places skill-intensive operations in advanced countries, labor-intensive operations in developing countries. As in the case of computers, there is clear horizontal specialization, with developing countries producing standardized commodity products – the manufacture of "standard chips," which are used in many devices, is dominated by emerging Asia, but much higher-end production remains in advanced countries.

There is also extensive vertical specialization. For example, Intel's manufacturing facilities are of two kinds, because production takes place in two stages. First, circuits are printed, using photolithography, on large disks of silicon at "wafer fabs". Then the wafers are sent to assembly and test facilities, where, according to Intel's fact sheet, "each wafer is cut into individual silicon dies, placed within external packages, and tested for functionality." [4]

Where are these operations located? Intel has wafer fabs in the United States, Ireland, and Israel. All of its assembly and test sites, by contrast, are in developing countries: they are located in China, Costa Rica, Malaysia, and the Philippines, with a new site under construction in Vietnam. In other words, within microprocessor manufacture, which is just one piece of the 6-digit semiconductor sector, one stage of production is largely confined to advanced economies, while another is largely confined to very low-wage countries.

All of this indicates that data showing a rapid rise in developing country exports, and Chinese exports in particular, within sectors that are skill-intensive in the United States need to be taken with large doses of salt. As Jin (2006) puts it, "The kind of gap seen in the electronic information industry between the rapid expansion of the scale of the industry coupled with a low value-added structure is evidence for China's role as an assembly base that is dependent upon overseas parts, intermediary goods, and capital goods."

Schott (2007) offers additional evidence based on unit values, which are available in sectors where the goods the U.S. imports have natural physical units. (Schott: "Examples of the units employed to classify products include dozens of shirts in apparel, square meters of carpet in textiles and pounds of folic acid in chemicals.") It turns out that Chinese goods imported by the

[4] See the Intel fact sheet: http://www.intel.com/pressroom/kits/manufacturing/manufacturing_qa.htm

United States have substantially lower unit values than goods within the same industries imported from OECD countries – e.g., the shirts we import from China are cheaper than shirts imported from advanced countries. Furthermore, the gap in unit values has been rising over time, suggesting that the relative sophistication of Chinese exports within any given industry has been declining.

The broad picture, then, is that the apparent sophistication of imports from developing countries is in large part a statistical illusion. Developing countries in general, and China in particular, are probably specialized in very different niches within industries than the United States. But how does all of the bear on the question of whether rising trade with developing countries has led to rising wage inequality in the United States?

Several recent analyses, notably Schott (2004) and Lawrence (2008) have argued that such specialization in effect protects advanced-country workers from distributional effects of trade by placing OECD countries in a different "cone of diversification" from developing countries. Figure 14 (skip Figure 13 for now) illustrates the concept of cones of specialization, using a Lerner diagram. The curves X, Z, and Y represent unit-value isoquants – combinations of skilled and unskilled labor input that produce an equal value's worth (say $100) of each good at current market prices. The downward-sloping lines NN and SS reflect relative factor prices in advanced countries and developing countries, respectively. E_N and E_S show the aggregate factor endowment of each country.

As drawn, the figure is consistent with a pattern of specialization in which both advanced and developing countries produce Z, advanced countries also produce skill-intensive X, and developing countries also produces labor-intensive Y. Because both types of countries produce Z, the two factor-price lines represent equal value; in advanced countries, the cost of producing X is the same as that of producing an equal value of Z, so both goods can be produced there; the same is true of Z and Y in developing countries; but producing a unit of X is more expensive in developing countries than in advanced countries, while producing a unit of Z is less expensive. Each country is able to fully employ all its workers because its endowment lies in the cone of diversification illustrated in Figure 14.

The famous proposition that trade leads to equalization of factor prices – a proposition closely linked to the Stolper-Samuelson effect -- applies only if countries lie in the same cone. So the suggestion that developing and advanced countries lie in different cones may seem to obviate concerns about the distributional effects of trade. Thus Schott (2004) asserts that

"If all countries produce all goods, unskilled workers in the U.S. can be adversely affected by a drop in the world price of labor-intensive products …Specialization, however, means that U.S. firms produce a capital-intensive mix of goods and are therefore not threatened by cheap imports."

Lawrence (2008) makes a similar argument. And in fact Krugman (1995) suggested that the prospect of specialization offered one reason to doubt whether the distributional effects of trade could get much larger than they were in the early 1990s.

But the evidence on specialization within industries, and vertical specialization in particular, calls this interpretation into doubt. The shock behind rapid growth in developing-country exports of manufactured goods does not appear to be developing-country growth leading to falling prices of traditional exports such as apparel. Instead, what we seem to be looking at is a breakup of the value chain that allows developing countries to take over labor-intensive portions of skill-intensive industries. And this process can have consequences that closely resemble the Stolper-Samuelson effect.

This point is difficult to make analytically, but comes across clearly in a numerical example. [*Note: in revision I will attempt to calibrate the example to bear a "stylized" resemblance to actual data*] For this example I assume that there are two final goods, Y and Z, produced using two factors of production, skilled labor S and unskilled labor L. There are also two countries, skill-abundant North and labor-abundant South. Production of Z is unskilled-labor-intensive. Production of Y takes place in two stages: a skill-intensive "component" stage X and a labor-intensive "assembly" stage.

Production functions and utility for the final goods are assumed to be Cobb-Douglas. Table 6 shows the assumed parameters and resource endowments.

The model is initially solved for equilibrium in the advanced country North in the absence of trade. I then consider two trade scenarios. Case I involves the assumption that X and Y must be physically co-located, so that there is in effect an aggregate XY industry. Case II allows X and Y to be separated, with labor-intensive assembly in a different country from skill-intensive component production.

Solution of the model requires determining both the pattern of specialization and relative goods prices. In practice I guessed at the specialization pattern, solved through a process of successive approximation for the equilibrium prices of both X and Y in terms of Z given the assumed pattern, then checked to confirm that the implied factor prices did in fact support that pattern.

In the case in which X and Y had to be co-located, the pattern of specialization and the associated factor prices were as illustrated in Figure 13. The developing country specialized in the production of labor-intensive Z, while the advanced country remained unspecialized, producing both the skill-intensive composite XY and Z. The relative price of Z was lower in the advanced economy than it would have been in the absence of trade, so there is a standard Stolper-Samuelson effect. As the first column in Table 7 shows, trade raises the real wages of skilled workers while reducing those of unskilled workers.

But what happens if X and Y can be separated? Then the pattern of specialization becomes that illustrated in Figure 14. Both countries continue to produce Z; meanwhile, the labor-intensive portion of XY moves to the developing country while the skill-intensive portion remains in the advanced country.

Figure 15 schematically illustrates the pattern of trade associated with each case, with the length of the arrows indicating the value of exports from each country to the other. When X and Y must be co-located, North exports Y to South and imports Z. (Think of this as trading computers for apparel.) When it becomes possible to engage in vertical specialization, North exports X – e.g., computer components – to South, and imports both Z (apparel) and Y (assembled computers). I've drawn the figure to suggest a large increase in the volume of trade. In the numerical example, the share of imports in North's GDP rises from .136 to .535.

In a qualitative sense the change illustrated in Figure 15 seems to resemble the actual change in North-South trade since the early 1990s, as documented in this paper. The share of advanced country GDP spent on imports from developing countries rises sharply, because components are shipped to developing countries for assembly, and the assembled goods are then exported back to the first world. If X and Y continued to be classified as part of the same industry, however, factor-content calculations based on advanced-country unit input coefficients would not suggest an increase in effective imports of unskilled labor. And the measured export mix of developing countries would seem to move upscale, toward more sophisticated products.

Yet as the second column in Table 7 shows, the actual effects on workers in the advanced economy would reflect a sort of Stolper-Samuelson effect: the real wages of skilled workers would rise, while those of unskilled workers would fall. Intuitively, the new ability to outsource labor-intensive industry segments to the third world would depress the demand for less-skilled workers, a shock not captured by data that lumped labor-intensive “assembly” operations together with skill-intensive “component” manufacture.

We should note, however, that this example does suggest that the type of calculation performed by Bivens (2007), in which the distributional effects of trade are assumed to be essentially proportional to the import share – a calculation suggested, admittedly, by my own 1995 paper – may exaggerate the distributional effects of recent trade growth. In this example, the trade share grows fourfold, but the distributional effects do not grow in proportion. [*Again, this example will be revised to become more realistic in the next draft*.] The reason, intuitively, is that much of the content of the new imports from developing countries is actually skill-intensive production from advanced countries [*in the current version of the example, 90 percent*], so that not as much unskilled labor is displaced as the raw import figures seem to suggest. If the United States imports computers from China, and China assembles computers largely from components made in Japan, only the assembly share of the sales price reflects labor-intensive imports – the rest is indirect importing from a country whose factor prices are similar to our own.

Nonetheless, the analysis presented here indicates that the rapid rise in manufactures imports from developing countries probably is, indeed, a force for growing inequality, and that factor content calculations suggesting otherwise are missing the essence of what is happening.

Implications of the analysis

The starting point of this paper was the observation that the consensus that trade has only modest effects on inequality rests on relatively old data – that there has been a dramatic increase in manufactured imports from developing countries since the early 1990s. And it is probably true that this increase has been a force for greater inequality in the United States and other advanced countries.

What really comes through from the analysis here, however, is the extent to which the changing nature of world trade has outpaced our ability to engage in secure quantitative analysis—even though this paper sets to one side the growth in service outsourcing, which has created so much anxiety in recent years. Plain old trade in physical goods has become remarkably exotic.

In particular, the surge in developing-country exports of manufactures involves a peculiar concentration on apparently sophisticated products, which seems at first to put worries about distributional effects to rest. Yet there is good reason to believe that the apparent sophistication of developing country exports is, in reality, largely a statistical illusion, created by the phenomenon of vertical specialization in a world of low trade costs.

How can we quantify the actual effect of rising trade on wages? The answer, given the current state of the data, is that we can't. As I've said, it's likely that the rapid growth of trade since the early 1990s has had significant distributional effects. To put numbers to these effects, however, we need a much better understanding of the increasingly fine-grained nature of international specialization and trade.

REFERENCES

Bureau of Labor Statistics (2006), “International comparisons of hourly compensation costs for production workers in manufacturing,” www.bls.gov.

Bivens, J. (2007), “Globalization, American wages, and inequality,” EPI Working Paper 279.

Blinder, Alan (2006). *Offshoring: The Next Industrial Revolution*. Foreign Affairs Magazine.

Cline, William (1997). Trade and Income Distribution. Institute for International Economics.

Dean, J., Fung, K.C., and Wang, Z. (2007), “Measuring the vertical specialization in Chinese trade,” USITC Working Paper.

Faberman, R. (2004), “Gross job flows over the last two business cycles,” BLS working paper 372.

Hummels, David; Ishii, Jun; and Yi, Kei-Mu. “The Nature and Growth of Vertical Specialization in World Trade.” Journal of International Economics 54 (2001): 75-96.
38

Krugman, Paul (1995), “*Growing World Trade: Causes and Consequences*”, Brookings Papers on Economic Activity, Volume 1.

Lawrence, Robert (2008), *Blue-collar blues: Is trade to blame for rising US income inequality?*, Institute for International Economics.

Mishel, L, and J. Bernstein and S. Allegretto (2006), *The State of Working America*, Cornell University Press and the Economic Policy Institute (EPI).

Romalis, J. (2004), "Factor proportions and the structure of commodity trade," American Economic Review.

Schott, P. (2004), "One size fits all? Heckscher-Ohlin specialization in international production," NBER WP #8244

Schott, P. (2007), "The relative sophistication of Chinese exports," NBER WP #12173

Stolper, W. and Samuelson, P. (1941), "Protection and real wages," Review of Economic Studies.

Table 1: Well-known estimates of the effect of trade on wages

Study	Estimated effect on skilled-unskilled wage ratio	Date of data
Krugman (1995)	3%	1992
Lawrence (1996)	3%	1993
Cline (1997)	7%	1993
Borjas, Freeman Katz (1997)	1.4%	1995

Table 2: Hourly compensation in top 10 U.S. trading partners as % of U.S. level

1975	1990	2005
Canada	Canada	Canada
Japan	Japan	Mexico
Germany	Mexico	China
UK	Germany	Japan
Mexico	UK	Germany
France	Taiwan	United Kingdom
Italy	Korea	Korea
Brazil	France	Taiwan
Neth	Italy	France
Belgium	China	Malaysia
76	81	65

Source: Compensation levels from BLS(2007). Trade values from Statistical Abstract. Note: Chinese compensation level for 1990 set at 1% of U.S.. Malaysian level in 2005 estimated from TK.

Table 3: Growth in GDP and manufactured exports

	1990	2006
China:		
GDP as % of US	6.7	20.0
Manufactured exports % of US GDP	0.24	2.13
Mexico		
GDP as % of US	4.6	6.4
Manufactured exports % of US GDP	0.37	1.16

Table 4: Workers “displaced” by trade

	1979-89	1989-2000	2000-04	employment, 2000
College graduate*	12.2%	21.2%	21.3%	25.6%
Non-college	87.8	78.9	78.7	74.4

Source: Mishel, Bernstein, and Allegretto (2007)

Table 5: 6-digit industries within Computer and Electronic Products

334413	Semiconductor & related device mfg
334111	Electronic computer mfg
334511	Search, detection, navigation, & guidance instrument mfg
334220	Radio & TV broadcasting & wireless communications equipment mfg
334210	Telephone apparatus mfg

Table 6: Assumptions for numerical example

Skilled labor share in X	0.8
Unskilled labor share in X	0.2
X share in Y	0.9
Skilled labor share in Y	0.01
Unskilled labor share in Y	0.09
Skilled labor share in Z	0.27
Unskilled labor share in Z	0.73
Share of Y in spending	0.5
Share of Z in spending	0.5
Skilled labor force in A	1
Unskilled labor force in A	1
Skilled labor force in C	0.1
Unskilled labor force in C	0.5

Table 7: Percent changes in real wages compared with no-trade situation

	XY co-located	Vertical specialization
Skilled workers	+13.4	+15.9
Unskilled workers	-11.8	-13.6
Import share in A	13.6	53.5

Figure 1

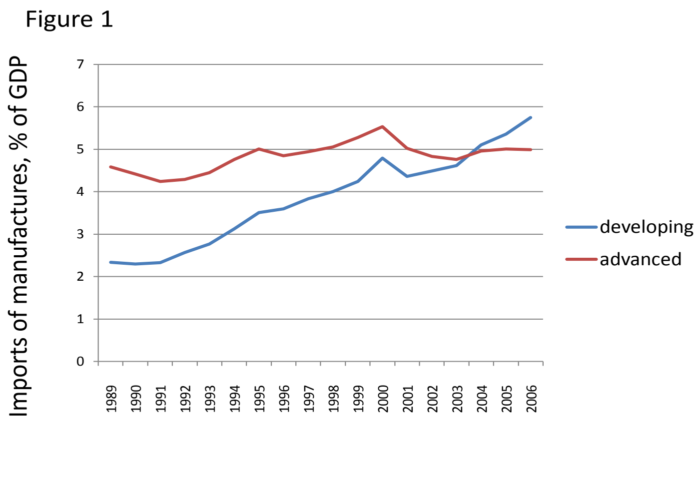
Figure 1
Imports of manufactures, % of GDP
7
6
5
4
3
2
1
0
1989
1990
1991
1992
1993
1994
1995
1996
1997
1998
1999
2000
2001
2002
2003
2004
2005
2006
developing
advanced

Figure 2

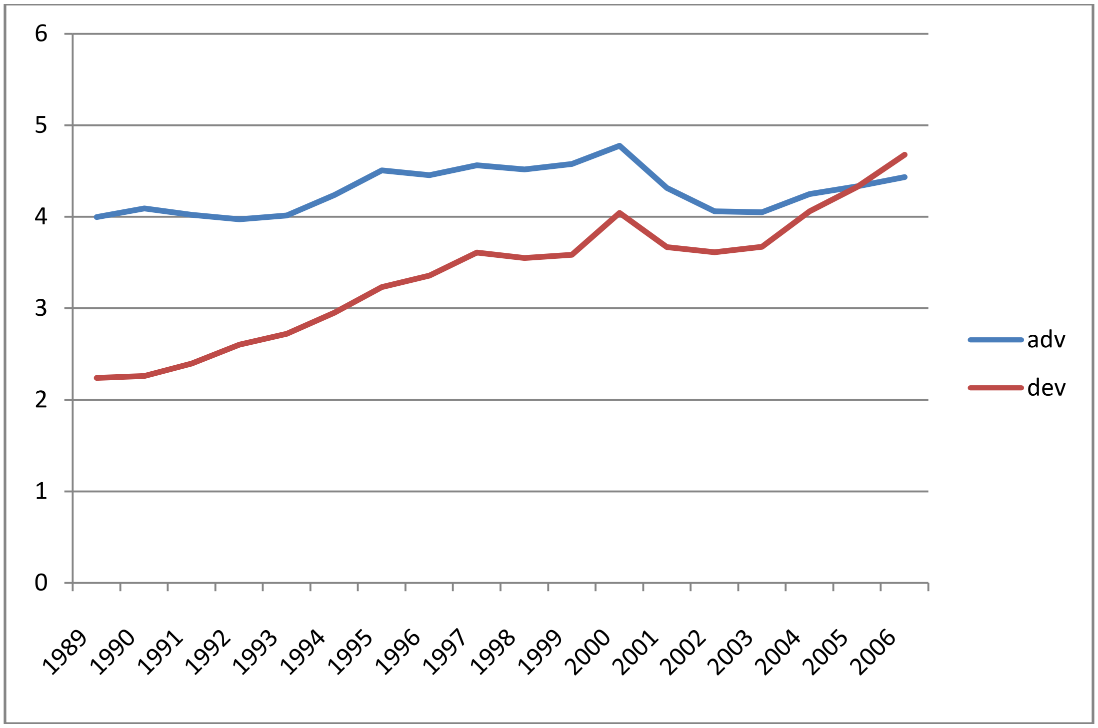
6
5
4
3
2
1
0
1989
1990
1991
1992
1993
1994
1995
1996
1997
1998
1999
2000
2001
2002
2003
2004
2005
2006
adv
dev

Figure 3

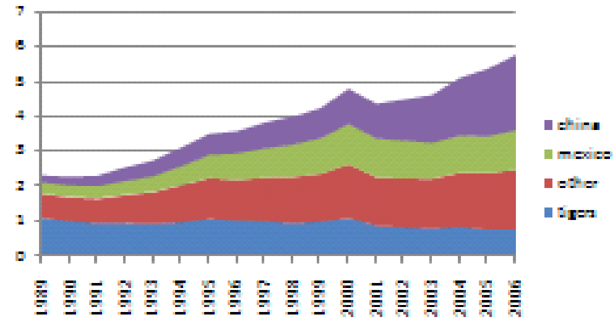

Figure 4

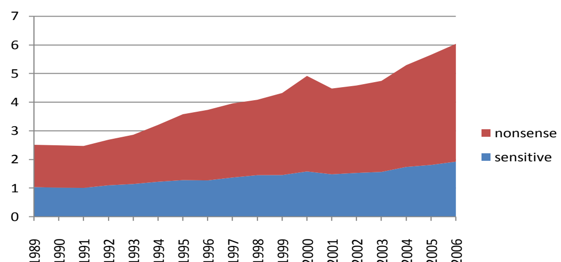

Figure 5

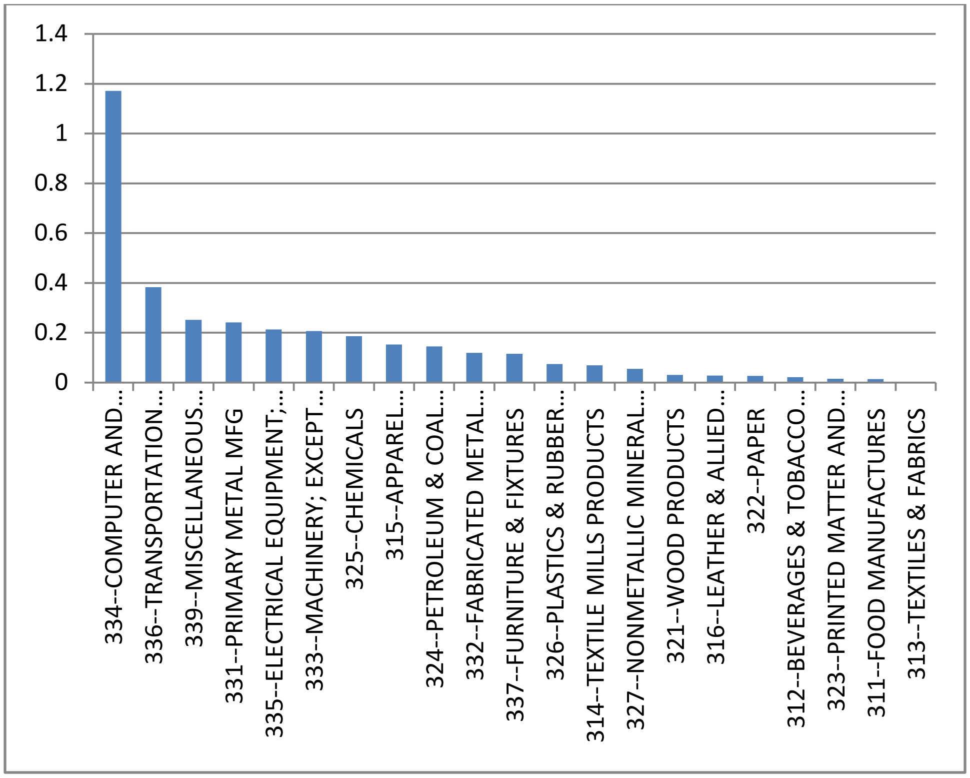

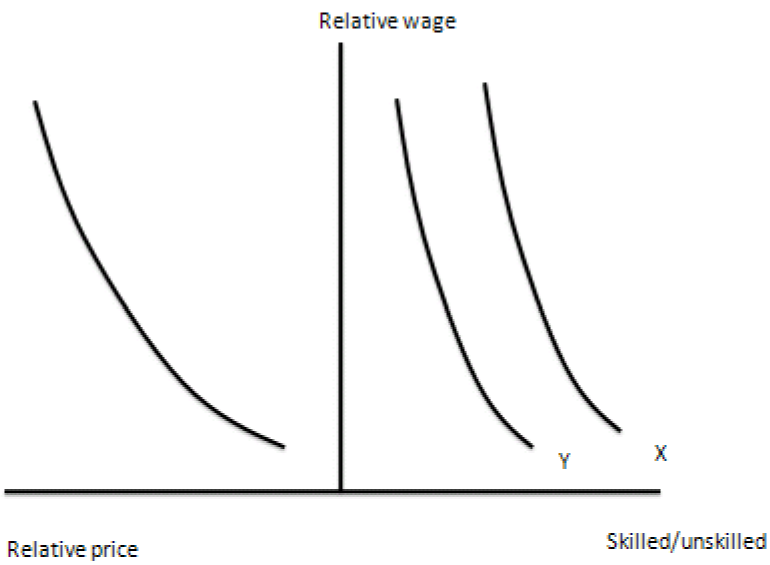

Figure 6

Figure 7

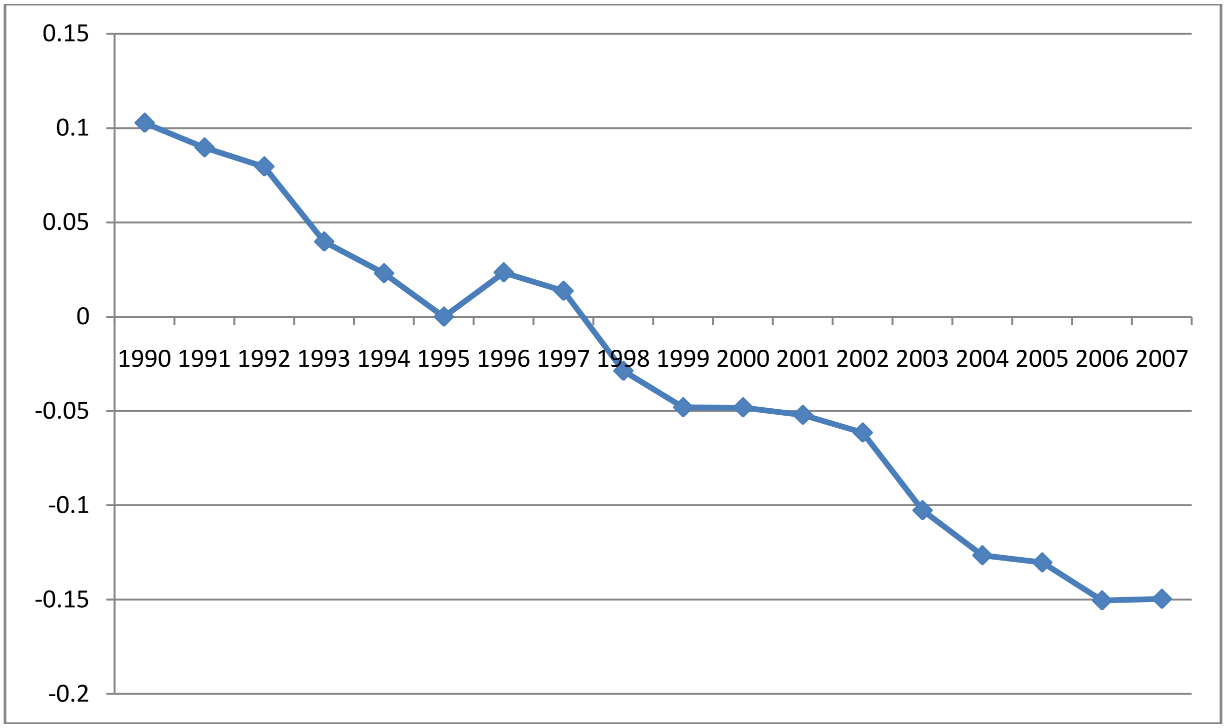

Figure 8

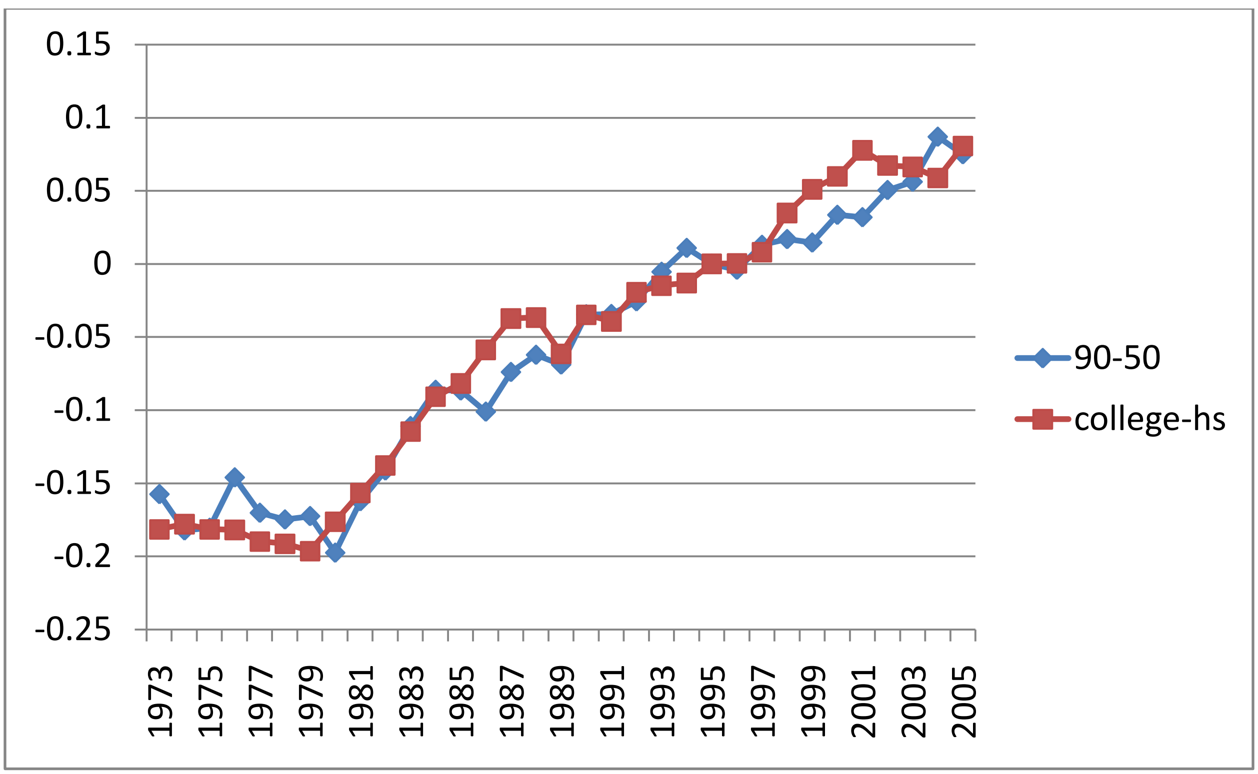

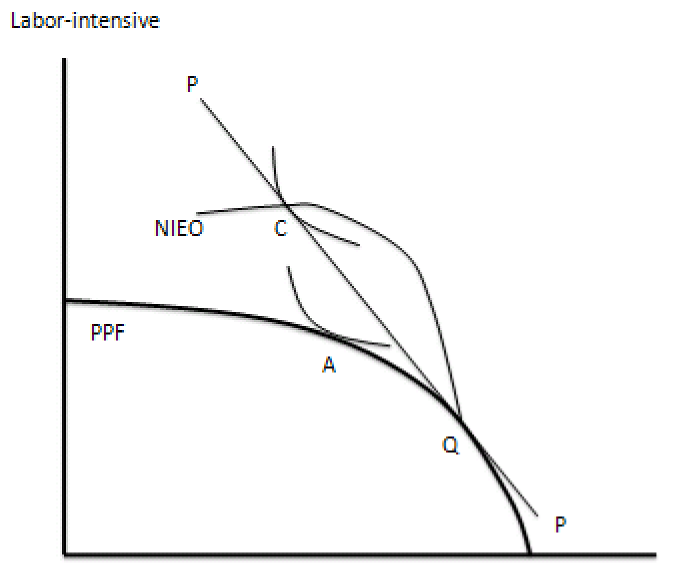

Figure 9

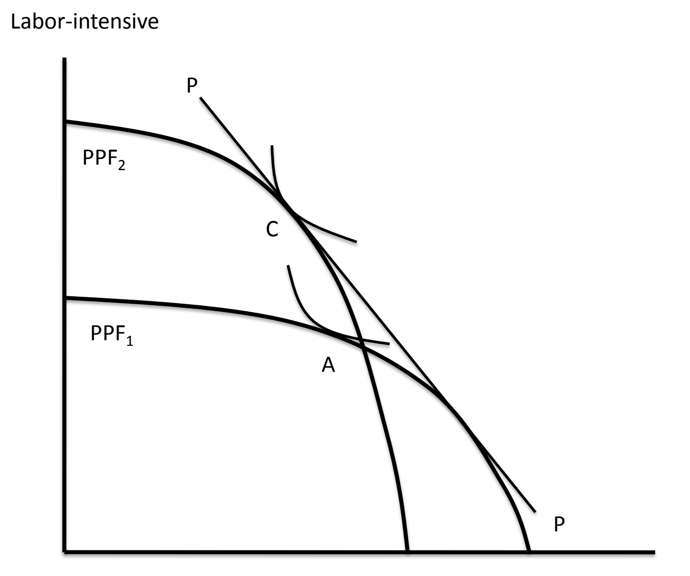

Figure 10

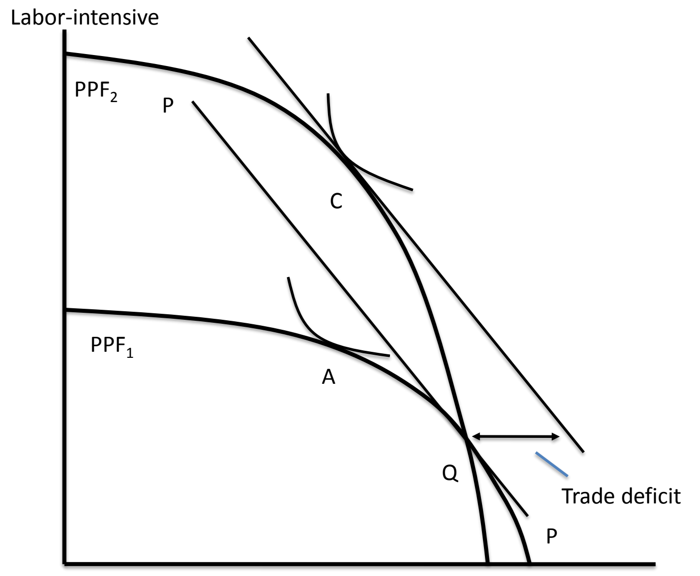

Figure 11

Figure 12

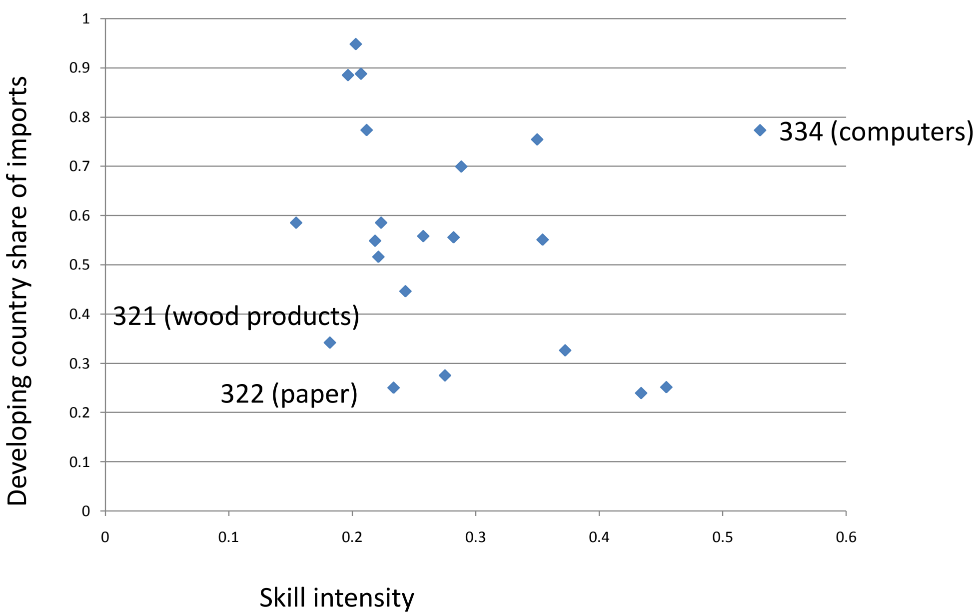

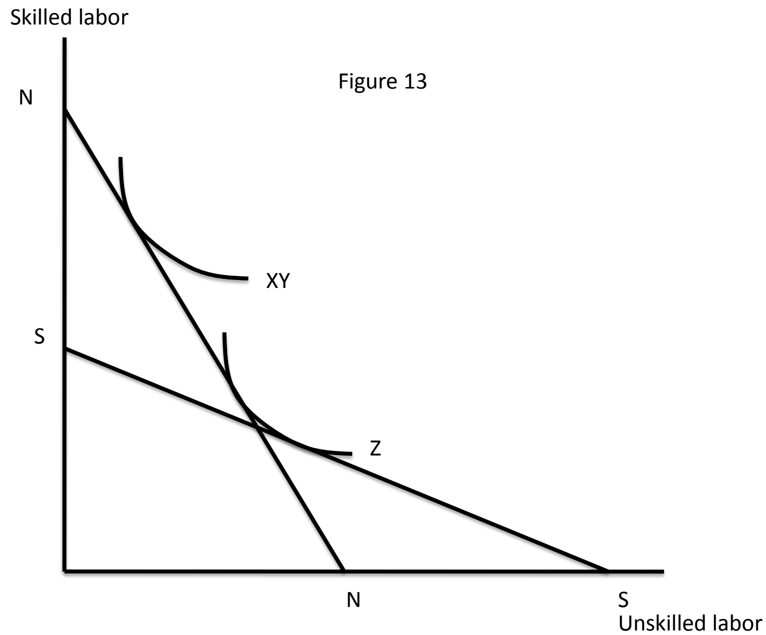

Skilled labor
Figure 13
N
XY
S
Z
N
S
Unskilled labor

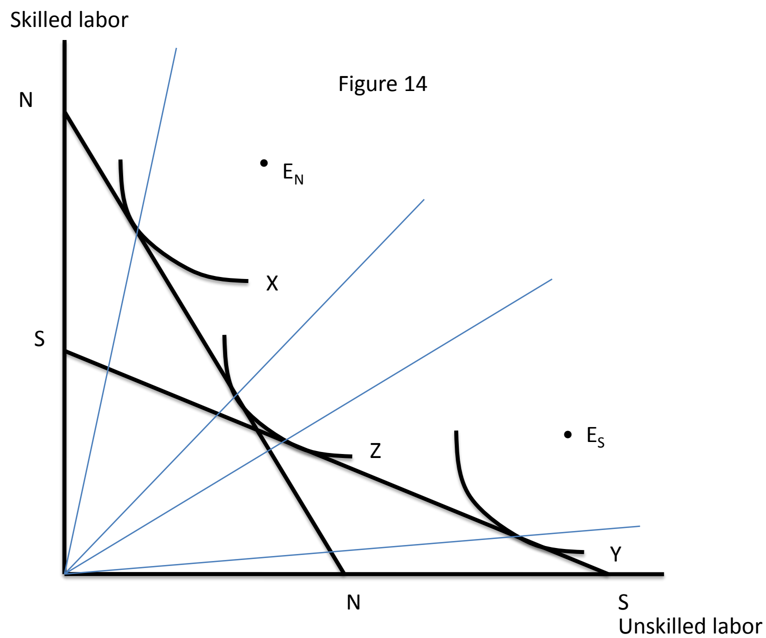

Skilled labor
Figure 14
N
E_N
X
S
Z
E_S
Y
N
S
Unskilled labor

Figure 15

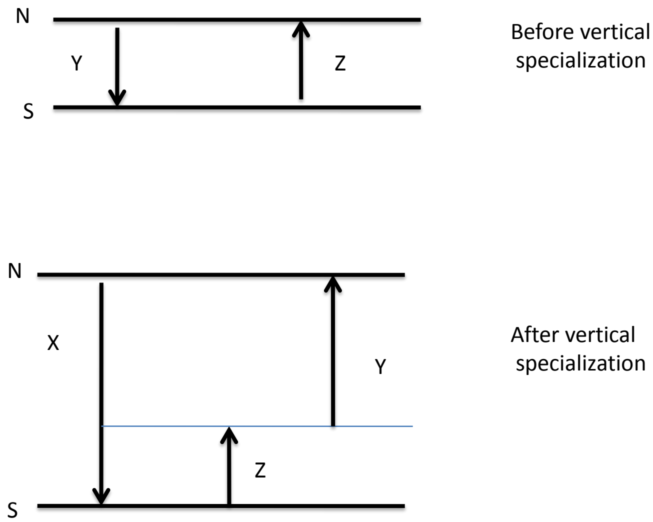
N
Y
Z
S
Before vertical specialization
N
X
Y
Z
S
After vertical specialization