

ON TRADE AND GROWTH IN A MARKET-ECONOMY SETTING

Arnold C. Harberger

University of California, Los Angeles

Paper presented at the Meeting of The Mont Pelerin Society

Salt Lake City, August 2004

We are living in a rather curious time with respect to trade policy, and to economic policy in general. There can be little doubt that the half-century between 1950 and 2000 was the greatest half-century in human history, in terms of the improvement of economic conditions and the betterment of life for the great majority of people. The last quarter of the 20th century was also probably the best quarter century of all time, from a strictly economic point of view. The Human Development Report of 2003 calculates the rate of per capita income growth for the developing countries at 2.3 percent per annum for the period 1975-2001. That compares with 2.1 percent for the high-income OECD countries. Both these figures are greater than the U.S. rate of per capita growth over the first half of the 20th century, when the U.S. was more clearly the leader of the growth race than it has been in the later period.

The picture for the 1990s was not much different, with developing country growth rising to 2.9 percent per capita per year and high-income OECD growth falling off slightly, to 1.7 percent per capita per year.

I think, with numbers like these we would be deafened by the cheers arising from people all over the world. Yet I have heard very little cheering, except from (and about) China and India. Instead one confronts a world full of complaints. In Latin America, there is a clamor against what they call neoliberalism and what most of us would call market-oriented economics.

In Washington, I have with my own ears heard serious statements (in high-level professional conferences) to the effect that the efforts at structural adjustment of the 1980s and 1990s had ended in failure, and so too had the so-called Washington consensus (which promoted policies of macroeconomic discipline, market orientation, more liberal domestic policies and greater openness to international trade).

How does international trade fit into this picture? Most importantly, the unprecedented world economic growth of the last half-century has been accompanied by an equally impressive growth of world trade. It is the exception, not the rule, to find a country or area whose international trade has not grown faster than its GDP. For most of the "growth miracle" cases -- Japan, Taiwan, Korea, Brazil, Spain, Portugal, Greece, Singapore, Hong Kong, Thailand, Malaysia, Indonesia, China, and now India, the miracle period was certainly characterized by a virtual explosion of exports, with exports growing much faster than GDP.

Many observers have applied the term "export-led growth" to these episodes, and I would not fault them for doing so, even though I would myself be more cautious, owing to my own appreciation of the complexity of the growth process. I feel there can be no doubt at all that the opening of these economies to freer international trade was an important factor in generating and supporting their respective "growth miracles". And, miracles aside, the importance of trade grew over the last 50 years in ordinary countries as well as miracle countries. One has no trouble in reaching the broad generalization that for just about every class of countries -- classified by stage of development, by form of government, by cultural or religious traits, by economic structure or by geographical region -- the relative importance of international trade has been growing over the last five decades.

Trying To Unravel The Paradox

To help us understand the paradox of widespread malaise and complaints at the close of a period of stellar economic achievement, I think the best approach is to go back to basics -- to see what economics has to say about international trade and about economic growth that might shed light on this paradox. This is what we will do next. But meantime let me anticipate the main conclusions, so that you know where I am going.

First, the main impact of trade liberalization is on the level of GDP, not on its rate of growth. The growth rate is affected as the economy transits from inefficient production (with restricted trade) to more efficient production (with freer trade), but once the transition to a more efficient economic structure has been completed, we have little reason to expect a major difference in the growth rate as a consequence of trade liberalization.

Second, the underlying mechanisms of economic growth lead their own lives -- bad policies can inhibit them and cause stagnation and misery, but good policies can't do much better than give them scope.

Third, what Schumpeter called "creative destruction" is an integral part of the growth process. The new displaces the old, which is happy news for those that are linked to the new and very unhappy news for those that are linked to the old. Real cost reduction is at the absolute center of the process of creative destruction. Most economists, analyzing the process, think of technical innovation as the key element. But the reduction of international prices is an equally effective way to reduce real costs for the users or consumers of these goods. Thinking along this line leads us to recognize how profound has been the revolution in low-end manufacturing during our lifetimes. The real costs of these goods have been driven down by the marriage of modern

production methods with sources of cheap labor. This has been a tremendous boon for consumers all over the world, but a corresponding bane to the “old” producers of these same low-end manufactures. We are seeing some of the same thing now in low-end services as they migrate to India, the Caribbean and other third-world destinations.

In sum, then, the prevailing malaise stems from: a) exaggerated expectations of the gains from freer trade; b) a failure to appreciate that economic growth comes mainly from inside the economy; government can help by letting it happen, but government can't make it happen; and c) unhappiness in places suffering the wrong end of creative destruction (which is fully understandable), combined with a quite general (and less understandable) failure to recognize it as an integral part of the growth process.

The Process Of Economic Growth

It was a bit more than 50 years ago that the economics profession hit upon a simple but profound insight -- that a country's (or industry's or business firm's) growth rate of output could be broken down into components. As in a recipe to make a cake, one adds more ingredients and one expects to get more products. But above and beyond the adding of ingredients, in economic production in many areas, we have been able to get more and more products from basically the same or equivalent ingredients. This comes from improvements in what most economists call total factor productivity. For myself, I prefer to call it real cost reduction, because that is a term most business people readily understand and even identify with.

A standard “breakdown of growth” will separate out at least three components -- one due to incremental labor occupied in the process in question; another due to incremental capital, and a third due to real cost reduction. A more sophisticated breakdown would also take into account

any changes in the skill mix of the labor force. Sometimes this leads to an estimate of growth due to incremental labor (assuming the skill mix stays the same) and another due to the change in the skill mix (change in average quality of labor). At other times people have broken the labor contribution down into a part due to raw labor (what would be the case if all the added labor were unskilled) and a part due to "human capital", whose contribution to product is measured, for workers of each type, by the excess of their wage over the unskilled wage.

Early work on the U.S. economy produced a breakdown of growth that went roughly as follows. Incremental labor accounted for growth of about $3/4$ of one percent per year. Incremental capital accounted for a like amount, and real cost reduction averaged about $1\ 1/2$ percent per annum. This breakdown is roughly descriptive of U.S. economic performance during the first half of the 20th century.

Since then there have been hundreds, probably thousands of different studies -- at the national level, the sector level, the industry level and even the firm level -- covering many different countries. The results of these studies are not easy to summarize, but I believe we are safe with the following generalizations.

- a) The labor contribution is typically less than the rate of growth of the labor force. The "constant quality" assumption leads to this contribution being the rate of growth of the labor force times the share of labor in GDP (usually $1/2$ to $2/3$). To exceed this figure, one needs improvement in average quality, which rarely adds more than half a point per year to the labor contribution. Rough numbers to keep in mind -- labor contributions of $1/2$ percent per year to 1 percent per year for advanced countries with slow-growing labor forces, and 1 to $1\ 1/2$ percent per year for developing countries with faster-growing labor forces.

- b) The capital contribution can be represented by the rate of net investment (as a fraction of a year's final product) multiplied by the rate of return attributed to that investment. (If you save 20% of your income and get a return of 6% on that amount, the resulting increment to your income is 1.2%). If we measure product by GDP or by value added (which is an industry's or a firm's contribution to GDP), the rate of return that we use should be gross of depreciation (reflecting the fact that gross domestic product is measured in this way). The rate of investment, however, should be net of depreciation, because it represents the net increment to the capital stock of the entity in question. Thus consider a country with a GDP of 100 and a capital stock of 300, with gross investment equal to 20, and with a 10% expected net real rate of return. If the relevant depreciation rate is 5% per year, net investment will be 5 ($= 20 \text{ minus } .05 \times 300$). And the contribution of net investment to growth will be 3/4 of one percent. This is the product of a rate of net investment of 5% and a gross-of-depreciation rate of return of 15%. This example represents a reasonable level for the capital contribution in most countries. Changing the rate of gross investment to 30% would raise the capital contribution to 2 1/4 percent per year. Changing the net rate of return to 15% would shift the result to 1.0 percent in the original case (gross investment = 20% of GDP) and to 3.0 percent in the modified case (gross investment = 30% of GDP).
- c) The real cost reduction term has been the wild card in the analysis of growth. Pick a set of cases of outstanding growth, and you will very likely find high rates of real cost reduction. Pick a set of cases of low or negative growth and you will very likely find low or negative rates of real cost reduction.

This is not the place to go into great detail on the real cost reduction story. But some discussion is needed because this element is so critical in understanding the growth process.

The first point to emphasize is that the real cost reduction that we measure all takes place, in one way or another, at the level of the firm. It takes a thousand different forms, but ultimately it responds to an urge in the part of business people (to reduce costs and thus add to profits) that is as basic and natural as the urge of consumers to maximize the satisfaction they get within the limitations provided by their income and wealth.

The second point to be noted is the frequency with which we observe negative real cost reduction. In large samples of firms and narrowly-classified industries, it is not at all unusual to find a third or more of annual observations to reflect real cost increases instead of reductions. And studies covering large numbers of countries can readily produce such cases as often as a quarter of the time. Why does this happen? The honest answer is we have not delved deep enough to be satisfied, but we have learned quite a lot. The main empirical regularity that we find is that cases of negative real cost reduction tend to appear simultaneously with declining output and declining rates of return. This in turn leads one to think of cases in which the entity in question faces a declining demand for its products. This can arise in two types of situations -- one, a cyclical or crisis-induced decline in demand, and two, reduction in demand for one's own products arising from intensified competition from other sources of supply (i.e., Schumpeter's creative destruction). Both of these elements are important. In our own work with country data, in a sample covering some 56 countries for periods of 20 to 40 years, we find nearly a quarter of annual observations to reflect real cost increases, with about a third of these occurring during cyclical declines and two thirds taking place in "normal" periods (neither cyclical declines nor

cyclical recoveries). We attribute these latter observations mainly to creative destruction, i.e., to competition from other sources of supply.

The final point to be noted is that three types of circumstance give a country a leg up in the drive to reduce real costs -- starting from way behind in the race, recovering from a cyclical slump, and reforming out of the morass created by bad economic policies. These circumstances can lead to observations of real cost reduction at rates of 4% per year or more. But our observations tell us that it probably takes rather special circumstances to reach such rates. The ordinary cases of real cost reduction tend to range between one and two percent per year, but with a lot of variability around this range.

The bottom line, then, is that it is cause for rejoicing when a country manages a growth rate of 3 or 4 percent per annum in real terms. And it typically takes rather special circumstances to exceed that rate. Nobody should consider it a failure for a country to be in this range. I am afraid that market economics has taken a bad rap from those who allege that the structural reform movement and the Washington consensus were failures. These people are asking too much. It is as if one were to develop a program to give special training to selected students from disadvantaged homes, and then say that the program failed because they didn't win National Merit or Rhodes Scholarships. In such a case, a B average in a good college or university should be enough to signify success. Similarly growth rates of 3 or 4 percent per year should be fully appreciated as representing good performances by a market-oriented economy. To work harder in such cases is always a good option, but to turn away from a market orientation because 3-4 percent growth is "not good enough" would be tragic and self-defeating.

Trade Liberalization and Economic Growth

The key point to be made on this topic is that freer trade permits an economy to make better use of its resources. This has been the message of market economics ever since the days of Adam Smith and David Ricardo. But to put it in a modern setting I know of no better example than an episode witnessed by Jim Henderson and me in Beijing in 1983. We had been sent there by the World Bank to conduct an intensive course in project evaluation. The participants in the course were employees of China's two main banks -- the China Construction Bank (which dealt with purely internal lending) and the China Investment Bank (which handled lending with international aspects). Our participants had had no training whatsoever in Western-style market economics, but they were very smart -- having passed through innumerable filters to get their bank jobs in the first place, and still further filters to gain placement in our course. The lectures in the course ran all day, with a break for a cafeteria-style lunch at which Jim and I mingled with the group. At one of these lunches, a couple of participants sat with us and told us their story of the automobile industry.

At that time almost the only cars to be seen on the streets of Beijing and other Chinese cities were Chinese versions of the 1942 Pontiac sedan, for which the dies and machinery had decades earlier been shipped to China. These cars weighed about two tons and had a voracious appetite for fuel. Sprinkled in among these behemoths, however, one could see a few contemporary Toyotas. The contrast was such that any visitor would notice it, and maybe it was Jim and I who brought the topic into the conversation. But anyway, the notable point of the conversation was our Chinese participants' reply. "You know," they said, "we started doing some calculations on autos. And what we found was that if we took the same value of resources

that was being used to make one of these big old cars, and shifted those resources to the textile and shoe industries, and if we then exported the textiles and shoes that these shifted resources produced in their new location, we could with the proceeds buy two brand new Toyotas". That's a pretty long sentence, but it's also a pretty efficient sentence, since it fully captures the essence of the principle of comparative advantage, which our two participants had rediscovered, all by themselves.

This is a good example because it shows how in this case, comparative advantage worked to create more value for the same resources. Trade liberalization works in the same way. If a country has a 50% uniform tariff that means that a dollar's worth of imports sells for, say 15 pesos while the dollar's worth of exports only brings the exporter 10 pesos. Reducing the tariff to 40% will stimulate trade, with the extra exports (costing 10 pesos per dollar) paying for extra imports valued at between 14 and 15 pesos per dollar. The excess of this benefit (14-15 pesos) over the cost (10 pesos) of the extra dollars represents the gain from the tariff reduction.

Using the same principle, consider a truly major trade liberalization, say from a uniform tariff of 50% to a uniform tariff at a 10% rate. Suppose, too, that this reform generates a quite spectacular increase in trade, with exports going from 10% to 30% of GDP, and with trade being balanced both before and after the change. The gain to the economy from such a major trade liberalization would amount to 6% of GDP.¹

¹This is obtained by considering that the "first" increment to trade has a cost of 10 pesos per dollar and a benefit of 15 pesos per dollar (reflecting the initial 50% tariff), while the "last" increment to trade has a cost of 10 pesos and a benefit of 11 pesos per dollar. The "average" benefit is thus equal to 30% $[(50\% + 10\%) \div 2]$. Applying this average benefit to the increment of exports (= 20% of GDP), we obtain 6% of GDP as the overall benefit of the liberalization.

Some people find it sobering, even disappointing, when they learn that the consequence of such a major liberalization is a benefit of “only” 6% of GDP. But they should realize that this benefit of 6% will go on and on into the future. If GDP were not to grow at all, its present value would be 60% of GDP at a 10% discount rate and 120% of GDP at a 5% discount rate. (The formula is present value = annual benefits ÷ discount rate).

If GDP is growing, the present value of the same 6% benefit gets bigger. At a 10% discount rate it amounts to 86% of the first year’s GDP, and at a 5% discount rate it goes to a whopping 300% of the first year’s GDP. (The formula here is present value = first year’s benefit ÷ [discount rate minus rate of growth of benefit]).

This is enough to show that generation after generation of economists were talking serious business when they pressed again and again for freer international trade. But note that these examples do not posit any permanent change in the growth rate as a result of liberalization. In the formula calculation above, the assumption was that right away, in the first year of liberalization, we see the full gain of 6% of a year’s GDP. That would mean that if the economy was growing at 3% normally, then for that one year there would be a growth rate of 9%, with the 3% growth rate prevailing from year two onward. This is grossly unrealistic, as the effects of trade liberalization come only gradually over time, as major resource reallocations occur -- toward export activities and away from those of import substitution in the previously protected sectors. Thus, the more likely scenario would be of 3% growth up to year one, and then of, say, 4% growth from year one through year six, followed by “normal” 3% growth thereafter. An alternative would be 3 1/2% growth from year one through year 12, again with a return to “normal” of 3% growth thereafter. Both of these scenarios provide a cumulative gain of 6

percentage points of GDP, with the gain here being spread over a transition period (of six or twelve years in the cases cited) rather than packed into just one year (as in the original example).

One way of expressing the thought that I am trying to convey is to refer to trade liberalization as affecting the level rather than the rate of growth of GDP. The rate of growth is not totally unaffected, but it changes only as a result of the transition from one level to another.

I cannot leave this topic without adding what I think of as a rather modest qualification. I earlier emphasized the important role that real cost reduction plays in the growth process -- how it is the factor that best discriminates between good and bad growth experiences, how it is a constant and never-ending objective of business people, how it is reflected in thousands of different ways, how it is very difficult to predict. With all of these complications serving as caveats, I think we can say that competition typically operates to stimulate real cost reduction. This effect stems from the fact that competition typically makes people work harder, strive more, put out more effort. Thus we can expect that in the more competitive situation that prevails after a trade liberalization, people in the affected industries will probably work harder to reduce real costs than they would have done under the umbrella of protection. There is modest evidence which suggests this is the case, but it is not yet conclusive. The reason is that the methods employed do not distinguish between the transitional effect on the growth rate and the more lasting effect to which this paragraph refers. I am currently planning a research program that will be precisely aimed at making such a distinction.

Meantime, we can be quite sure that simply having reasonably free trade does not automatically give a country a new and much higher growth rate. Enough countries have joined the freer trade group so that a large permanent effect on their growth rates should be readily

visible. I think we must therefore settle on a presumption that some permanent effect probably exists, but that it is only one of many factors accounting for observed rates of real cost reduction -- which, as noted earlier, have a combined central tendency in the range of 1 1/2 to 2 percent per year.

Growth As a Cumulation of "Changes in Level"

I feel a little reluctant to assert -- yet say it I nonetheless will -- that the idea of "self-sustained growth" can be seriously misleading. The reason is that most of the elements that produce growth do so by changing levels, not by impacting the rate of growth itself. We have seen how trade liberalization works mainly in this way. Let us now consider other elements in the growth process.

a) We add a new investment to the economy, which has, say, a 15% rate of gross return. What we may expect from this investment of, say, 1,000, is a flow of benefits of 150 at the start, declining by 5% per year as the investment depreciates. (This assumes a 10% rate of net return and a 5% per year depreciation rate). If the depreciation is straight line for 20 years the expected flow of benefits from the investment would start at 150 and drift down to 50 over the project's 20-year life. In general, each investment can be thought of as giving a positive pulse to the growth rate (+ 150 in these examples) followed by a series of small negative impacts as the contribution of this investment to GDP (measured by its gross-of-depreciation rate of return) declines. An extreme case would be an investment of the so-called "one-horse-shay" variety, which produces a constant service yield (S) throughout its life, and then collapses all at once. Here the impact would be a positive jump of +S period one, with zero contribution to growth thereafter, followed at the end of period N by a negative jump of -S.

b) We add to the education level of the labor force of the country, increasing its earning power by, say 12%. If labor's contribution to GDP was initially 500 out of a GDP of 1000, this change would raise it to 560, producing a 6% rise in GDP. I constructed this example so as to yield precisely the same increment to GDP as the trade liberalization discussed in the previous section. We get the same effect. In a zero growth setting the present value of the benefits of this change would, in the simplest case, be 600 at a discount rate of 10% and 1200 at a discount rate of 5%. If this change were superimposed on a path in which GDP was growing at 3%, its present value would be about 860 at a 10% discount rate, and would be all of 3000 at a 5% discount rate. Those calculations are based on a 6% extra jump of GDP in year 1, leading to a 9% growth in that year followed by no extra growth thereafter. The more likely scenario for such a change would be an extra 1/2 percent growth over 12 years, or an extra 1/4 percent growth over 24 years, as successive new cohorts of better-educated people joined the labor force.

c) A cost-reducing innovation is introduced, leading to an increase of real product from 1000 to 1060 using the same resources as before. Again the same figures as before apply. In a stagnation setting, where GDP is not growing, we have a steady stream of 1000 being converted to a steady stream of 1060. In a growth setting, we shift from a stream starting at 1000 and growing at 3% to a stream starting at 1060 and also growing at 3%. The growth rate leaps to 9% for a single transition year in the simplest case, followed by no growth effect thereafter. If the innovation is introduced gradually over an extended period (as was the case with hybrid corn, with antibiotics, with the assembly line, with computerization as well as with just about every other innovation) then the extra growth will be spread over time as in the earlier examples,

leading to an extra 1 point of growth over 6 years, an extra 1/2 point of growth over 12 years or an extra quarter point of growth over 24 years, as the case may be.

What lessons can we draw from these examples? The first lesson is that a “standard” impulse of growth, regardless of whether it impacts the capital contribution, the labor contribution or the real cost reduction term, will typically operate via a “level effect”, increasing the growth rate over a transition period, but not permanently.

The second lesson is that to raise the growth rate permanently, we must keep introducing new impulses to growth. We raise the average level of education from 8 to 9 years and we get growth at 3 1/2 percent instead of 3% for something like a decade. To keep it at 3 1/2 percent by this route we have to make additional efforts, bringing average education, say, from 9 to 10 years for the next decade and from 10 to 11 years for the one after that.

With physical capital, if we add one investment this year in a standard pattern, this has to be followed by another one next year and yet another the year after that. This is already captured in our way of representing the capital contribution to the growth rate, where the ratio of investment to GDP appears as one of the two components of this contribution. Adding to the rate of investment for just one year produces just a blip in the economy’s rate of growth. To have a permanent effect, we have to be shifting the rate of net investment from, say 10% of GDP to say, 12% of GDP. At a 15% gross rate of return, this permanent upward shift in the investment rate will change capital’s contribution from 1.5 percent per year to 1.8 percent per year -- i.e., it will add three-tenths of one percent per year to the country’s growth rate. But this means a different extra 2% of GDP being invested in different new investments in every successive year.

The 0.3 percent effect on the growth rate is the composite result of a perpetual chain of extra annual investments, each of which by itself is just contributing a “level effect”.

The bottom line of this section is that each given effort, regardless of whether it works to improve the labor contribution to the growth rate, or the capital contribution, or the contribution of real cost reduction, is likely to generate a permanent rise in the level of GDP, but will produce only a blip in the growth rate. Raising the growth rate requires a series of such blips, coming one after the other in rapid succession. And each of these blips requires its own push of extra effort. Extra growth does not come easily. It takes work -- lots of work, opening new vistas period after period. It is wrong -- often badly wrong -- to think that big improvements in the growth rate will come automatically, as the concept of self-sustaining growth appears to imply.

Now to a qualification of what I have just said. The exception occurs when a lot of bad, even stupid policies have been artificially holding down the forces of growth. Bad labor policies can deprive employers of the incentive to hire additional workers; bad education policies can end up producing only tiny increments of productivity and earnings as one moves up the educational ladder; widespread corruption and arbitrary interferences with economic processes can all lead to investments being made elsewhere, not here; rampant inflation can blunt people's perceptions of relative prices and costs, and major price distortions can cause people to miss genuine opportunities to reduce true real costs while pursuing false opportunities stemming from the wrong price signals. In short, every single component of a country's growth rate can be held down by wrong-headed policies. Such situations can yield negative or zero or miniscule growth rates for extended periods of time.

When a sensible set of reforms takes away these kinds of trammels to the efficient functioning of the economy, one can get first a significant spurt of growth as distortions are eliminated and as advantage is taken of long-neglected opportunities. But even after this phase of “recovery” or “emergence from the morass”, there can be a permanent effect on the country’s growth rate as it moves from negative or zero or miniscule up to a “normal” range of about 3 or 4 percent. This range reflects normal rates of investment, normal rates of return, normal growth of the labor force and its quality, and normal rates of success in the search for ways to reduce real costs. This kind of “normal” growth can legitimately be called self-sustaining, but it represents the functioning of natural forces. Government policies can and should open the door for these forces, but they cannot “create” them.

The Role of “Creative Destruction”

One can conveniently visualize the workings of creative destruction in three different settings -- that of standard commodities, that of differentiated products, and (in a sense overlapping the first two) that of international competition.

Hybrid corn provides us with an easy scenario for the first of these settings. As hybrids began to be introduced in the 1930s, some farmers were ready to take a chance and plant the new varieties. Those that were successful made a lot of money, and their neighbors and others proceeded to initiate them. As corn supplies grew, the relative price of corn fell, and the benefit of the innovation, which was in the first instance reflected in high profits for the hybrid-planting farmers, was in the end passed on to consumers in the form of lower relative prices and higher quality of corn.

In the beginning, before prices began to fall, farmers who planted the old varieties of corn were substantially unaffected. But as prices fell they found themselves squeezed. For the early adopters, the shift to hybrid corn meant higher profits (which then dropped gradually back to “normal” as prices fell). For the late adopters, it was a question of dealing with growing losses as prices fell; for them, shifting to hybrids became a matter of simple survival; in the end, they could not make it if they stuck with the old, traditional, but now inferior varieties of corn.

The above is the “big picture” of the hybrid corn revolution. But the little picture is equally germane. It so happens that given hybrid varieties were successful only in certain areas. The hybrid that worked in one county could well turn out to be a failure in the county next to it. As a result, there were many farmers who tried specific hybrids (which had done well in the experiment stations) only to find that they were ill-suited to local soil and weather conditions. These cases could easily result in reduced profits and in real cost increases for the affected farmers. Often, this type of experience developed as farmers tried to follow the lead of others in neighboring counties, only to find that those particular hybrids did not perform well in their own location.

So we had some negative real cost reduction in the cases of farmers who waited too long to convert to hybrids (and were then squeezed by falling prices); we had other cases of early experimenters who simply tried the wrong hybrid seed; and finally we had later failures as farmers tried to duplicate the good results of others, only to find that those varieties were not suited to their particular local conditions.

The case of differentiated products is well illustrated by the case of supermarkets squeezing out the old-fashioned “mom and pop” retail food stores, and by the way such chains as

WalMart, Target and K-Mart ended up squeezing out a lot of traditional department stores.

Sometimes different competitors will develop different technologies and one will lose out to another, as Sony's Betamax lost out in the VCR race. Sometimes it will be one business plan that wins out over another -- as IBM's strategy of licensing its technique to other manufacturers won out over Apple Computer's strategy of trying to "go it alone".

The main point here is that it is pretty hard to think of a major cost-reducing innovation that worked so evenly on all producers at the same time that there were no losers in the game. Losers are generally an integral part of the picture, and their losses typically give rise to real cost increases, which operate to partially offset the gains stemming from the real cost reductions of the winners. In a dynamic economy with rapid growth generated by lots of real cost reductions there are likely to be quite a lot of real cost increases suffered by competitors. The end result is typically either that these competitors go out of business, or alternatively that they end up by themselves following the innovators and adopting the innovation in question.

Special mention should be given to the case of foreign competition, if only because of the political overtones it carries. When the winners and the losers are both in the same country, the beneficence of an innovation is easier to defend. When the increased competition comes from abroad, an entire phalanx of resistance is very often formed by the domestic producers who are thus threatened; and the protectionist snake is once again poised to strike.

I think the world has been pretty lucky in recent decades to have been able to resist protectionist pressures as well as it did. We economists are well aware that the ultimate beneficiaries of real cost reduction are the world's consumers, and that in some deep sense the

benefits of an innovation that are enjoyed by consumers will in the end outweigh the costs that are borne by labor and capital in the activities that are threatened.

But I think it is wrong for us to be “Pollyanna freemarketeers”, trying to tell threatened textile workers or steelworkers that all will be well for them in the end. By far the better approach is to recognize their problems as real ones, but then to point out that many suggested solutions to these problems, particularly those involving protectionism in some form or other, nearly always carry economic costs that far exceed their benefits.

Competition from abroad, in the form of lower world prices for steel, shoes, textiles, whatever, is indeed a benefit for consumers of those products, no matter where they are located. This same competition is also a threat -- or, put more positively, a challenge -- to other producers of the same or competing products. As such it can generate real cost increases induced by declining demand, and can easily lead to a mediocre growth performance.

We are seeing something of this sort, I believe, as the world responds first to the challenges posed by the original Asian tigers, and now to the similar but even stronger challenges emerging from China and India as they become new major world centers for low-end manufactures as well as certain services. Life is tough for the sectors that compete with these activities, even in countries not specialized in them. But certainly life is tougher in countries that are more heavily specialized in these low-end products than in those that are lucky enough to have their comparative advantage in other activities. The booms now underway in China and India, for example, have helped to trigger rises in the relative prices of many primary products. They have helped the world's producers of oil and copper while making life quite difficult for producers of low-end manufactures. I regard all these effects as being almost “acts of God” in

the same sense as hurricanes, earthquakes, floods and wildfires are. By that I mean that they are things we all should accept as part of reality, and then figure out how best to respond.

One area in which we can all draw useful lessons from this analysis lies in sensitizing ourselves to these realities. We do not expect an outstanding growth performance from Honduras when half its banana trees have been blown down by a hurricane. Similarly we should not expect great performances from El Salvador and Mexico in periods when those countries' main products are beset by major competition from Indonesia, China and India. Finally, we should also understand that Chile's outstanding growth performance of recent years, which owes a great deal to a set of very sound economic policies, has also benefited from a booming demand for copper emanating largely from China and East Asia. If we maintain a careful and subtle appreciation of the circumstances of each country, we will end up being far better judges of the quality of its policy performance.