

Thinking Liberal

thinkingliberal.co.uk

Matthew Green - essays on economics



2. Trade & globalisation

ONE OF THE REASONS I chose to study economics was the discipline's capacity to deliver counter-intuitive insights. You hear somebody spouting away with ideas that seem to be common sense – but economics gives you the insight to see deep logical flaws. Nowhere is this more true than in the subject of trade. We've all heard the arguments that low rates of pay in China will drive down living standards here unless globalisation and competition is restrained. The insights that show this logic to be deeply flawed are hardly new: they were first publicised by Adam Smith and David Ricardo two hundred years ago. These insights are highly relevant to the current globalisation debate, but too much energy is lost arguing over the basic concepts, and not enough spent considering the full implications of the globalisation phenomenon. Here's the punch line: the explosion of physical world trade that we see now is not a permanent change to the way the world is: it is part of a transition from an unequal world to one where wealth is more widely distributed. As that transition completes, trade will reduce, and the benefits to the rich world will reverse..

Let me take you on a journey. I will start like countless economics lecturers over many generations by looking at two farms, side by side. One farm is in the valley, where the soil is good and access easy. The other is in the hills, with poor soil and steep slopes. Whatever type of farming you want to do, the cost per unit of output is more favourable in the valley farm. Does that mean that the valley farm out-competes the hill farm and puts it out of business? The answer is no, unless the valley farm is able to satisfy all demand for all products by itself. The valley farm has a limit to its production; once it's met its capacity then the hill farm has no competition for any extra goods.

Now let's extend this. The two farms are on an island; each farm supports its own community and between the two they cover the whole island. Trade with the outside world is difficult. Let's start with the proposition that each farm and community is self-sufficient. The hill farm community is poorer, and the valley farm richer because of the relative productivity of their farms. But apart from relative productivity, there is an important difference between the farms, and this is critical to trade theory: the relative productivity of different farming products is different. The valley farmer faces a trade-off, or opportunity cost, between wheat and sheep, say, of 1:1. For the hill farm this ratio is 1:10. So for each sheep each farm takes out of production, the valley farmer can replace it by growing ten times the wheat than the hill farmer would. The valley farm may still be more efficient for sheep farming, but it is over ten times more efficient at growing wheat: the hill farm's land is completely unsuitable for growing wheat, but it is tolerably productive for sheep. So provided both farms want both wheat and sheep products (wool, meat, cheese) in reasonable quantities they can gain from trade. The hill farm stops growing wheat, and buys it all from the valley farm. The valley farm stops farming sheep: it buys all of its sheep products from the hill farm. This is a more efficient use of their

combined resources and both sides get richer. The hill farm can release a lot of space from its attempts to be self-sufficient in wheat: more than enough to satisfy their neighbour's demands, since there is a 10:1 trade off. But although both sides gain, the valley farm is still richer than the hill farm.

This is the basic logic of trade theory and it's easy enough to grasp if you take it slowly. In fact we come across it often. Say Jill is a high-flying executive, and Jack is a humble chauffeur; Jill is better at housework and looking after the kids than Jack, but Jack is still perfectly capable. It makes sense for Jill to go out to work and Jack to stay at home, even though Jill knows she could do a better job at home. She can't do both. Economists call this comparative advantage. The valley farm has a comparative advantage in wheat, and Jill in employment; the hill farm has a comparative advantage in sheep and Jack in housework. Everybody wins if they specialise in the areas where they have comparative advantage.

Unfortunately the expression "comparative advantage" is not well understood in the world at large, and it doesn't really do what it says on the tin. It seems to say "advantage compared to" and sounds a lot like "competitive advantage". But the hill farm's comparative advantage in sheep doesn't fit this: the valley farm can farm sheep more cheaply, and if it chose to compete with the hill farm on sheep, it would beat it. But it can make more money by using its resources to farm wheat, so it doesn't try. Probably the key insight is that resources are limited. If our two farms had limitless land and/or labour, they wouldn't bother to trade. But they have to choose between devoting resources to sheep or wheat.

And while we're on our island, it's worth pointing out that something else doesn't make sense. It doesn't make sense for the hill farm to stop production and for its community to live entirely off the valley farm. Why would the valley farm let them do this? What does it get in return? Of course if the hill farm could send workers to work on the valley farm, that picture could change. So what's the point of China destroying British industry and supplying us with all our needs with a massive one-way trade surplus? We couldn't pay for it. It would be an act of charity.

So let's draw the wider lessons. If the opportunity costs of producing different, tradable goods varies between two communities, and transport costs are acceptable, there will be advantages to trade to both sides. Each community has a comparative advantage in at least one good, and this is the good it specialises in and exports. But the trade must balance out if both sides are to advantage; if they don't balance then it is the community with the export surplus that loses out. That community might as well take its export surplus to the docks and push it into the sea. In the long run. These are the insights that Adam Smith and David Ricardo gave us two hundred years ago, and it's not difficult to see that this clashes with much conventional wisdom. Lower costs don't automatically mean that there is a trading advantage; export surpluses aren't automatically a good thing.

But this theory doesn't explain how comparative advantage between large, diverse entities like countries arises. And frankly, economists have struggled with this: there have been lots of ideas (especially an idea called "factor endowments") but at high level at least, these don't do a good job of explaining many current trading patterns. But one area where the theory seems to work very well is trade between countries at different stages of economic development – where comparative advantage arises from different rates of technology development. This is what we will look at next.

I will do this by conducting a thought experiment using a simple economic model. This will offer some striking insights. Let's start by considering a country with 100 population units – perhaps think of them as millions of people. The economy is divided into four sectors. First is agriculture: no development story can be understood without it. To keep matters simple we will assume that agricultural products are not tradable; each country wants to be self-sufficient, and cannot export its surplus. Next we have services. We will take a narrow definition of what services are: essentially demand for people's time that must be supplied within a community; think of schooling, perhaps. By definition this cannot be traded between countries; neither is it possible to improve productivity – this builds on the idea developed in my first essay. And finally we have two sectors of tradable goods: high-tech and low-tech. As a country develops, its productivity in high-tech goods advances faster than low-tech. In practice this is likely to be because educational standards increase, which has a bigger effect on high-tech production than elsewhere. At the start of the development journey we define units of output to be what one population unit produces in one period. So all sectors start with a productivity of one. Now what do people consume? First they must meet a basic need for agricultural produce of 0.75 units per population unit. With their remaining resources people choose to acquire units of output in the following ratio: 5% on further food, 5% on services, 45% each on high and low tech goods. This is based on units of output, not cost and share of budget. The numbers are arbitrary but this pattern of consumption is broadly consistent with the sweep of history. We next assume that there is fully functioning and competitive labour market, and that all people are paid the same. This implies that everybody has the same basic skills. This is breathtakingly unrealistic, but it makes the picture a lot easier to see: changing the assumption would not change the shape of the conclusions. We next define a currency unit as what each population unit is paid in a time period at the start of development. The cost consumers pay for goods is the cost of making it. So all goods cost one unit per unit at the start. As time progresses each currency unit continues to buy one unit of produce on average, but the price of each type of product will change because the relative cost of producing it changes.

Now let's consider the effect of development. We will do this in two giant leaps. In the first step, agriculture is revolutionised and productivity rises to ten. Low-tech productivity rises more modestly to five, and high-tech advances to ten. Services stay at one. In the second leap, agricultural productivity advances from ten to fifteen; low-tech doubles from five to ten, but high-tech quadruples, advancing to 40. Now let's pause to see what effect this has had on the whole economy.

Stage	<i>Consumption (units)</i>			<i>Production (population)</i>			<i>Unit cost</i>		
	0	1	2	0	1	2	0	1	2
Agriculture	76.3	99.3	118.3	76.3	9.9	7.9	1.00	0.56	0.63
Services	1.3	24.3	43.3	1.3	24.3	43.3	1.00	5.62	9.42
Low tech	11.3	219.1	390.1	11.3	43.8	39.0	1.00	1.12	0.94
High tech	11.3	219.1	390.1	11.3	21.9	9.8	1.00	0.56	0.24
Total	100.0	561.8	941.8	100.0	100.0	100.0			

This shows the logic of industrial development quite strikingly. Overall consumption, and hence living standards, rise dramatically. The agricultural sector (look at the production columns, showing where people are employed) collapses both because productivity rises and because demand relative to other goods falls. Services rises, from 1% to 43% of the employment, because demand rises with general consumption, but productivity doesn't. This is known by economists as the Baumol effect, after the person that first identified it. It is often considered a Bad Thing ("Baumol's disease"), because it implies that lack of growth in productivity of services is a brake on total productivity in the economy – but this was the subject area of my first essay.

Now let's consider the effect of trade. Think about two economies, Dollarland and Yuanland. Dollarland starts our story at Stage 2, and stays there. Yuanland starts at the beginning, and works its way through the stages.

At the stage zero for Yuanland, there is a clear basis for trade, since the opportunity costs between high- and low-tech goods is different (compare 1:1 to 4:1). First we must set an exchange rate for the two currencies (dollars and yuan, of course); for this the residents of Yuanland must find Dollarland's high-tech goods cheaper than their own, while Dollarland's residents must find Yuanland's low-tech goods cheaper. An exchange rate of two yuan to the dollar does this, and we'll use this since it makes the maths easier. The dollar price of all of Yuanland's products is \$0.50, compared to Dollarland's costs of \$0.94 and \$0.24 for its low tech and high tech goods respectively. As we open up trade, Yuanlanders import high tech goods madly until their demand is fully satisfied and there is no domestic production left. Dollarland imports low tech goods until there is an overall trade balance. Workers are redeployed across the economy to more efficient activities and both economies get a dividend. The total value of trade is \$3.5 (7% of Yuanland's economy, 0.4% of Dollarland's). Yuanland gains 7.8%, Dollarland by 0.3% - it is a much bigger economy. The distribution of benefits is a function of the exchange rate – and the convenient 2:1 rate favours Yuanland.

The whole outcome seems to fit quite well with our view of the world economy. Yuanland is very backward. Its workers are paid just \$0.50, compared to \$9.42 in Dollarland. A visitor from Dollarland to Yuanland would find services absurdly cheap: \$0.50 compared to \$9.42. The market exchange rate multiplies the differences between the countries, just as we observe in the current world (purchasing power parity would imply a 1:1 exchange rate). Both countries are better off, but it is easy to see the tension. Yuanland's dollar economy is small compared to Dollarland, and its position looks weak: it needs the trade more than Dollarland does. Further, in open trade its high-tech sector cannot exist, which may limit possibilities for future development. Meanwhile in Dollarland people from the low-tech sector will complain that it is unfair to compete with people just paid just 5% of domestic wages; they will say that trade serves to perpetuate starvation wages – when in fact the low wages of Yuanland arise from its massive and unproductive agricultural sector. That this simple model predicts these commonly observed effects shows how powerful it is.

Now let's move the story on. Yuanland pursues a successful path of economic development and reaches Stage One on the path, while Dollarland stays where it is. The opportunity cost difference between low- and high-tech goods for Yuanland narrows to 2:1, but is still different to Dollarland's 4:1. The exchange rate (taking the nearest round number again) narrows to 1.5:1. Yuanland's now much higher income drives a big increase in trade, to over \$56 (6% of Dollarland's economy, 14% of Yuanland's). Yuanland's economy gains by 8.3%, Dollarland's by 1.5%. Yuanland's demand for hi-tech goods is still the limiting factor, and its hi-tech sector remains wiped out. The differences between the countries have narrowed. Yuanland's dollar

wage is now 40% of Dollarland's. This pattern of trade again seems to fit our experience of trade between developed and middle-income countries, like China perhaps, where volumes of trade become very significant to both sides.

So what happens next? Yuanland's development relentlessly catches Dollarland's and reaches Stage 2. This is where it gets interesting. As the opportunity costs become more similar, the incentives to trade, and the gains from trade, decline. Once the catch up is complete, under the comparative advantage theory, there is no trade at all. This is bad news for Dollarland, since it loses its gains from trade without any compensating advantage. For Yuanland the loss of trade gains is compensated by the advance in technology. This striking result was first publicised by the eminent economist Paul Samuelson in 2004. This was given two types of paranoid spin: first that it showed that conventional trade theory didn't always produce gains (in fact this is a case of gains being reversed); second it showed that as developing economies advance, they will threaten developed countries' export industries. But just as importantly, the developed countries would start finding that imports were getting more expensive.

But wait. If very similar economies don't trade because there is no comparative advantage, how come trade volumes are so high between countries like France, Britain, Germany and Italy? In fact the theory of comparative advantage doesn't do a good job of explaining trade between advanced economies that are geographically close together. But the volumes of trade between advanced countries that are far apart is proportionately quite low – as between the US and Europe. Perhaps trade with Japan is an exception, until we remember that Japan routinely runs large export surpluses with the developed world, and doesn't import much from them – so we are not explaining direct, two-way trade. There is every reason to think that trade with China and India will decline as these countries become more advanced. They will become more interested in producing goods for domestic consumption than for export.

And this is the idea that needs to be better appreciated. China won't hollow out Western economies leaving us unemployed or paid starvation wages. In fact what we'll find is that Chinese imports will start to become more expensive; the domestic industries that they rendered uncompetitive could start becoming competitive again. And yes, they will be increasingly difficult to export to, and be big competitors on world markets – but then since we will need to import less, we will need exports less, so this isn't a big deal. But we will find that a relentless increase in import prices will erode living standards. In 2008, in the Britain and the US at least, it looks as if this process has already started.

This essay was written in 2008, with only a minimal edit in 2016; the financial crash was gathering speed, but had not yet impinged on most people's consciousness. It was intended as the first of a series, but I ended up writing just one more.

*Matthew Green
February 2016*