MIT Open Access Articles

Comparative advantage and agricultural trade

The MIT Faculty has made this article openly available. ***Please share*** how this access benefits you. Your story matters.

Citation	Donaldson, Dave. 2019. "Comparative advantage and agricultural trade." Agricultural Economics, 50 (S1).
As Published	http://dx.doi.org/10.1111/agec.12529
Publisher	Wiley
Version	Author's final manuscript
Citable link	https://hdl.handle.net/1721.1/140492
Terms of Use	Creative Commons Attribution-Noncommercial-Share Alike
Detailed Terms	http://creativecommons.org/licenses/by-nc-sa/4.0/

Massachusetts Institute of Technology

DSpace@MIT

Check for updates

Comparative advantage and agricultural trade[1]

Dave Donaldson

Department of Economics, Massachusetts Institute of Technology

Received 10 June 2019; received in revised form 11 October 2019; accepted 17 October 2019.

Abstract

This article is based on a transcript of the Leonard K Elmhirst Memorial Address, delivered at the 30th International Conference of Agricultural Economists in Vancouver, Canada, on 29 July 2018. The presentation provided an overview of work that I have done recently on comparative advantage, its role in agricultural trade patterns and trade-driven changes in welfare. Two topics are addressed: (i) the static gains from intra-national agricultural trade due to comparative advantage; and (ii) how trade might mitigate consumption losses and food insecurity from changes in agricultural productivity caused by drought and climate change.

Key words: Comparative advantage; international trade; intra-national trade; agriculture

JEL codes: F15; N57; N77

---

[1] Corresponding author. E-mail address: ddonald@mit.edu. Elmhirst Lecture presented at the 30th International Conference of Agricultural Economists, Vancouver, Canada, July 28 – August 2, 2018. Current address: 77 Massachusetts Ave., Cambridge, MA 02139

This is the author manuscript accepted for publication and has undergone full peer review but has not been through the copyediting, typesetting, pagination and proofreading process, which may lead to differences between this version and the Version of Record. Please cite this article as doi: 10.1111/agec.12529.

This article is protected by copyright. All rights reserved.

Author Manuscript

Thank you very much for the chance to be here—it's a great pleasure and a real honor. I remember vividly in grad school when I discovered some of the early handbooks of agricultural economics in the library, and felt that this was the missing textbook in "How to be an applied economist". It was those sorts of lessons that pushed me in the direction of some of the work I ended up doing in graduate school and afterwards. And so I'm obviously delighted to be here to tell you about some of that work, as well as to have the chance to learn more over the next few days.

Will (Martin, President of the IAAE) asked me to speak about some of the recent work I've been doing. This is under the theme of comparative advantage and its role in agricultural trade patterns and trade-driven changes in welfare. This obviously draws on work that I've done with lots of other people. I'll be referring to work in particular with Robin Burgess (e.g. Burgess and Donaldson, 2010), and Arnaud Costinot and Corey Smith (e.g. Costinot and Donaldson, 2016; and Costinot, Donaldson and Smith, 2016). I am a trade economist at heart. In fact, I view the area of "trade" fairly liberally—so I think most economists are trade economists: we're studying trade between one group of people and another group of people, often on two sides of a market or possibly on two sides of the world. The questions that come up in this area – the policy questions – the ones that get me, and many of us, out of bed in the morning, are very hard. They typically involve a bewildering number of interacting components in a general equilibrium framework, with very little experimental or quasi-experimental variation on which we can draw as empirical scientists. And so we lean heavily on theory to cut through the fog. I view that theory as helping us to map the raw things we see in the data to the questions that require more than just what is in the data.

This article is protected by copyright. All rights reserved.

Author Manuscript

I am going to talk today about some theory, going back to Ricardo's now more-than-200-year-old body of work. It's the model of comparative advantage, but it's really the simplest one you can start thinking about, and I'm going to talk about two aspects of the lessons from Ricardo, using some data.

My talk today will cover two broad themes: The first is about levels. If you live in an economy that is opening up to trade, the "levels" question is: how large are the static gains from that new opportunity to trade? In the second part I'll be talking about volatility, which of course relates to the issues of food security that Will Martin just mentioned in his Presidential Address. The question here is: can trade mitigate volatility? So that's the roadmap for today.

1. Trade and productivity

I am going to start—maybe foolishly, but Will encouraged me, as my old physics teacher used to say, never to underestimate how much people enjoy hearing what they already know—with Ricardo's famous example, right out of an Econ 101 textbook. This is Sunday morning. It's early. So we'll start with an easy one. Here is a simple table with what has been called Ricardo's "four magic numbers" (Samuelson, 1972, p. 378). In Ricardo's example there were two countries and two goods. You could call them regions or people, and sectors or activities, or whatever else you want to call them. But let's stick with two countries, England and Portugal, and two goods, which in the famous example

Author Manuscript

This article is protected by copyright. All rights reserved.

were cloth and wine. Ricardo's four numbers are replicated in Table 1. Each hypothetical number in this example represents the number of workers that will be required in that country to make a unit, however defined, of the good in question.

Table 1: An illustration of comparative advantage

	Cloth	Wine
England	100	120
Portugal	90	80

Note: Cells represent the number of workers it takes to make one unit of the good

Source: Eaton and Kortum (2012)

As you can see in this example, it takes 100 English workers to make a unit of cloth, but it would take 120 such workers to make a unit of wine. In Portugal things are different, and differences are at the heart of comparative advantage. England has a comparative advantage in producing cloth and Portugal a comparative advantage in making wine. Another way of illustrating this is with the familiar production possibilities frontier (PPF) for two countries (Figure 1). These have different slopes, which is the key difference that gives rise to comparative advantage.

Another way of seeing this that's slightly less common in the Econ 101 textbooks is illustrated in Figure 2. This is the world production possibilities set. The dots show possible patterns of production specialization in this world economy. To the top left, there's a dot that represents how much world

Author Manuscript

This article is protected by copyright. All rights reserved.

output would be achieved if both countries produced only wine, and to the bottom-right if both produced only cloth. The dot to the north east, the upper right, is what happens if England makes the world's cloth and Portugal the wine. The not-so-happy point that's inside the PPF is what would happen if we got things backwards and England made the world's wine and Portugal made the made the world's cloth—just imagine! One can see right there the logic of comparative advantage driving gains from trade. In a trading equilibrium, from the global perspective any of the points on the PPF are possible production bundles for the world as a whole, and, hence, they're also possible consumption bundles for the world as a whole, and points interior to them are inferior. Anything that involves countries not trading is going to involve a bit more of a mixture of countries doing the wrong thing. Exactly how this world splits up those gains from trade involves further argument, but we usually think that both sides will enjoy at least some of it, so there's a strong argument there for gains from trade for both parties.

Author Manuscript

This article is protected by copyright. All rights reserved.

Author Manuscript

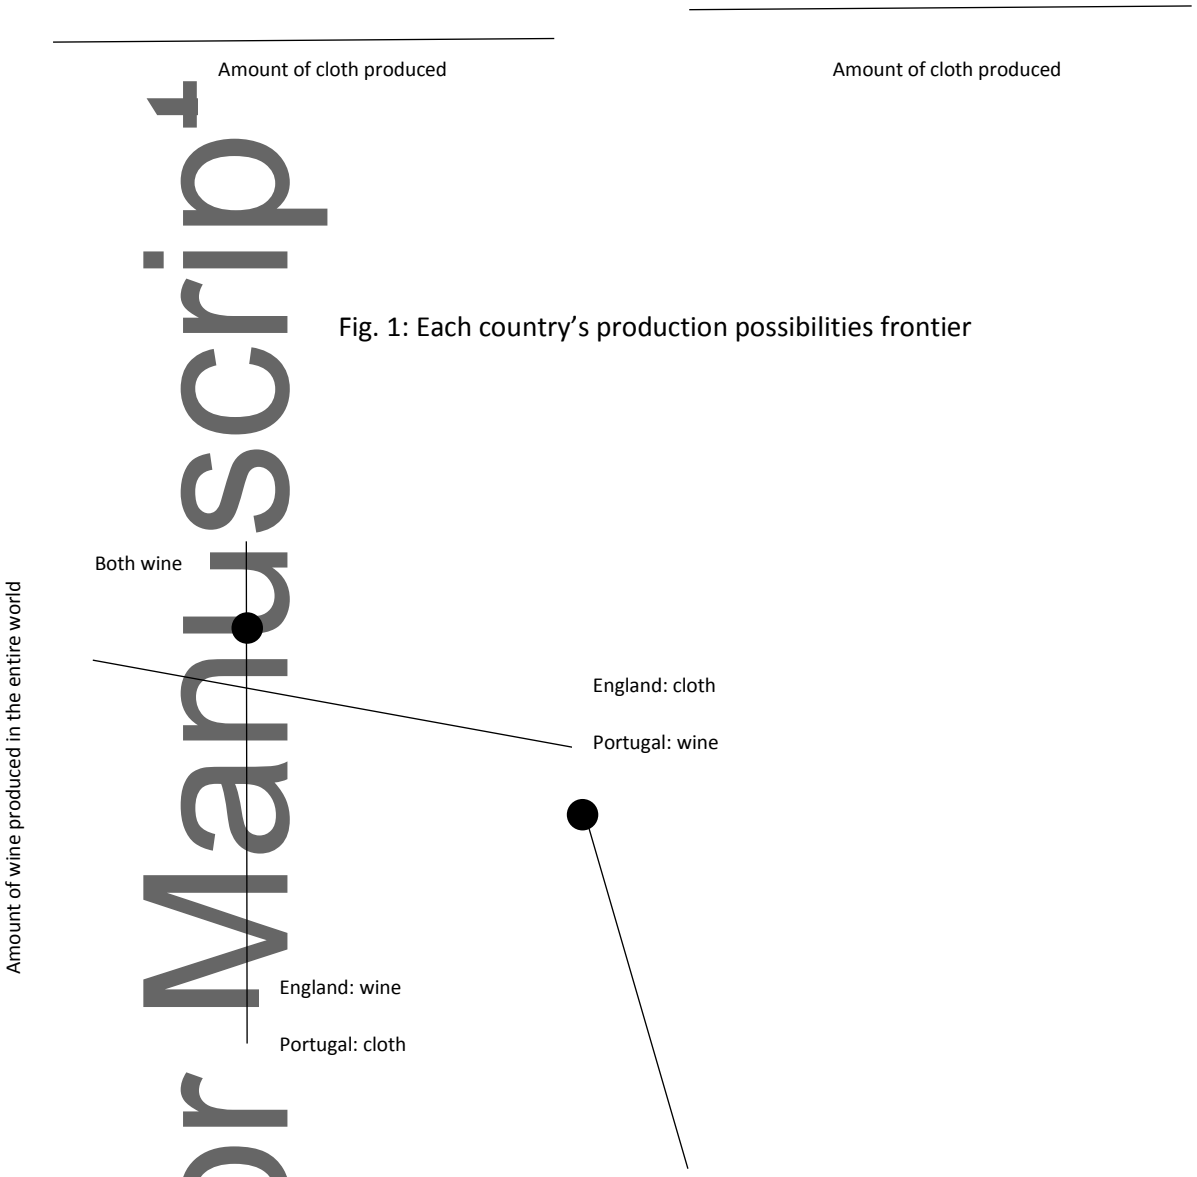

Fig. 1: Each country's production possibilities frontier

This article is protected by copyright. All rights reserved.

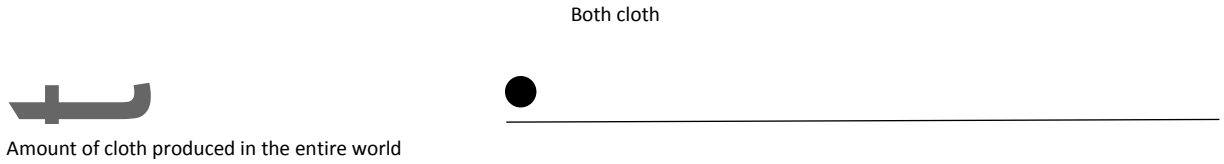

Fig. 2: Global PPF and comparative advantage for two countries

This simple example can be generalized slightly by adding a third country. There are still two goods, but now I'm adding France with some hypothetical productivities in cloth and wine, as reflected in the world PPF in Figure 3.

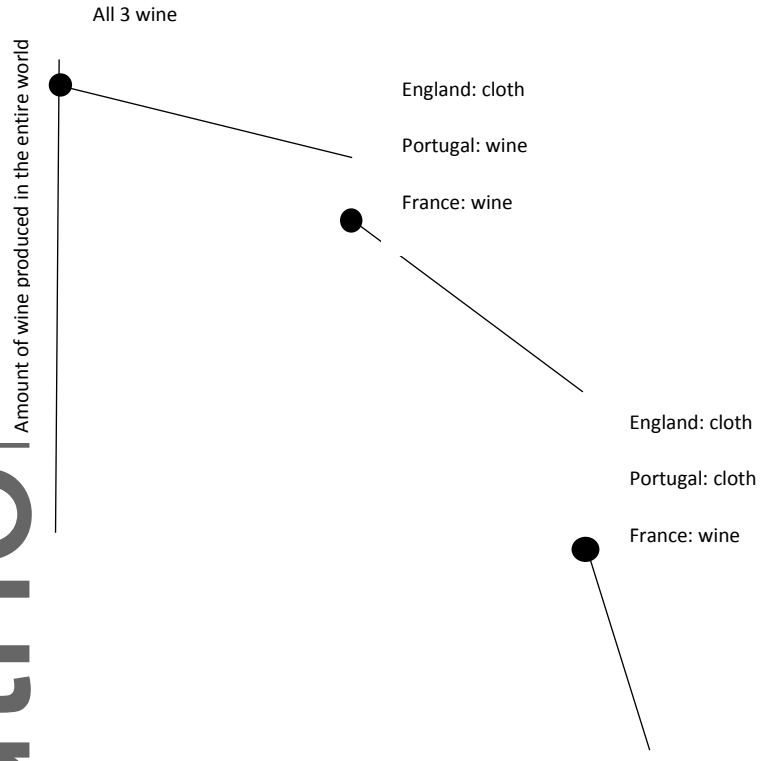

Author Manuscript

This article is protected by copyright. All rights reserved.

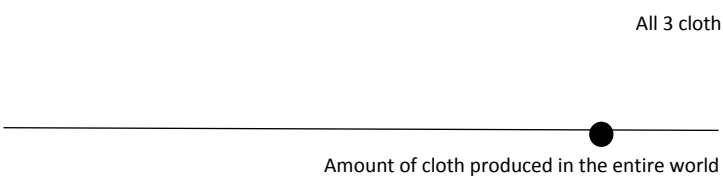

Fig. 3: Global PPF and comparative advantage for two countries

This looks similar to Figure 2, with just an additional, third flat segment. We can generalize that point and imagine a world with many such small units of analysis. Think of them as parcels of land or sub-countries somewhere in the world, for example, and the PPF now starts to look relatively smooth. Of course, in the limit, it would be perfectly smooth as we go to the extreme of many infinitesimally small countries.

Another way of seeing the gains from trade is illustrated in a picture that I'll refer back to throughout this talk (see Figure 4). In this, I've partitioned the world into two countries, but each country has a number of different parcels of land so each country has a smooth but otherwise well-behaved PPF. In this example, the PPFs are crossing. If these countries trade freely and without any transport costs there would be a common relative price. For a wide range of different types of preferences the slope of the dotted line would be this unique relative price. The countries can trade and maintain balanced trade (that is, without the need to borrow from anybody) and hence allow consumption anywhere on that dotted line. And that's the beauty of trade: you can produce at the dot on your PPF, but consume elsewhere on the dotted line. Another way of saying this is that additional consumption possibilities are opened up for either country (illustrated by the area under the dotted line, but above the PPF).

Author Manuscript

This article is protected by copyright. All rights reserved.

Author Manuscript

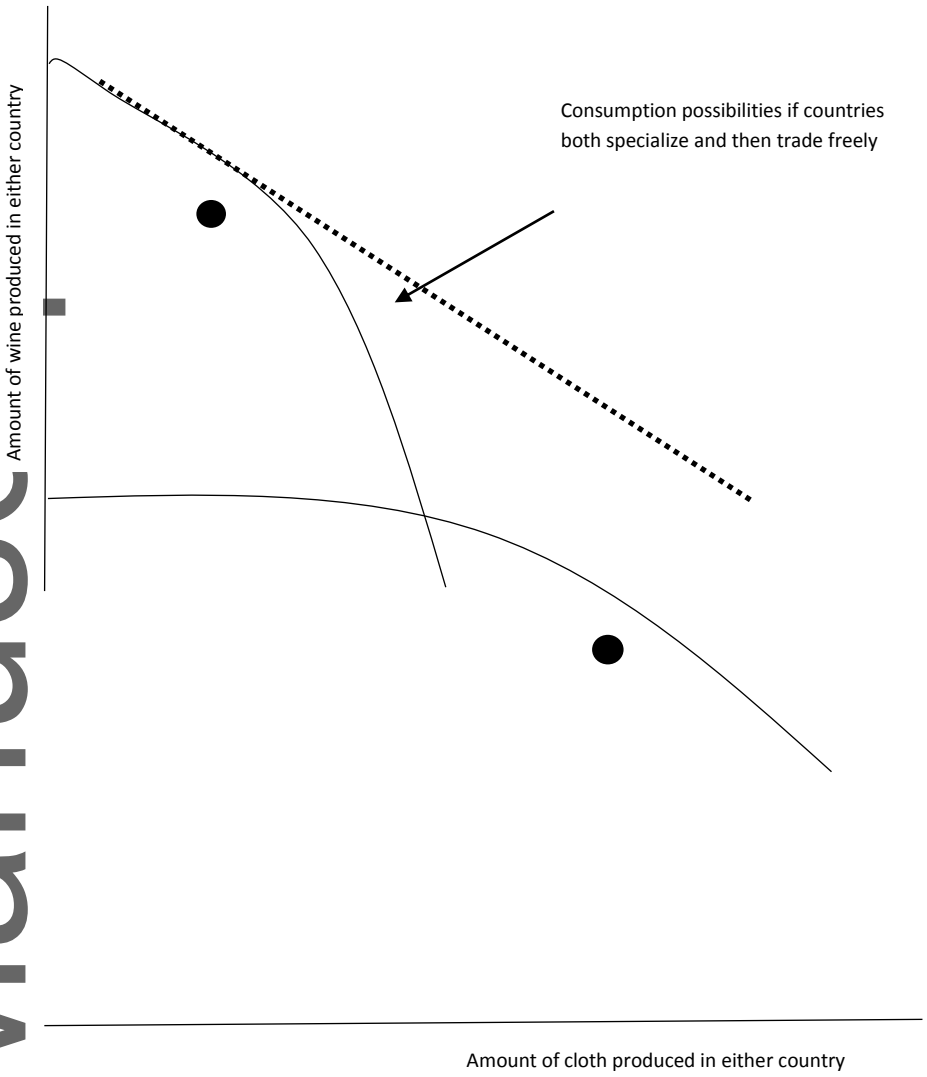

Fig. 4: PPFs of 2 countries, each with many sub-regions

This article is protected by copyright. All rights reserved.

Now consider Figure 5. What I'm trying to illustrate here is a possibility where you, the analyst, know the PPFs in your region, and you, early in time, see alternative production bundles (as represented by the smaller dots). Imagine that, over time, each country's smaller dots move in the direction of the arrows out towards the large ones on the dotted line. What could be going on here?

One possibility would be that, early on, these countries are not trading. So they might be facing barriers to trade that mean the prices that prevail inside country one or country two are different. That would be one rationalization for why we might see production taking place at these smaller dots that are away from larger dots (where there is one common world price and hence no price distortion). But if we were to observe these two countries with their small production bundle dots, and saw that these small dots were moving closer and closer to the dotted line, then we would see that there are aggregate gains being realized even as each country remains on its PPF. In that sense, there are no within-country productivity gains, as neither country's own PPF is changing, but the world is becoming more allocatively efficient—more productive as a whole—as prices move closer and closer to being equalized. Another way of saying the same thing is that if there are lower distortions, there are gains from trade.

Author Manuscript

This article is protected by copyright. All rights reserved.

Author Manuscript

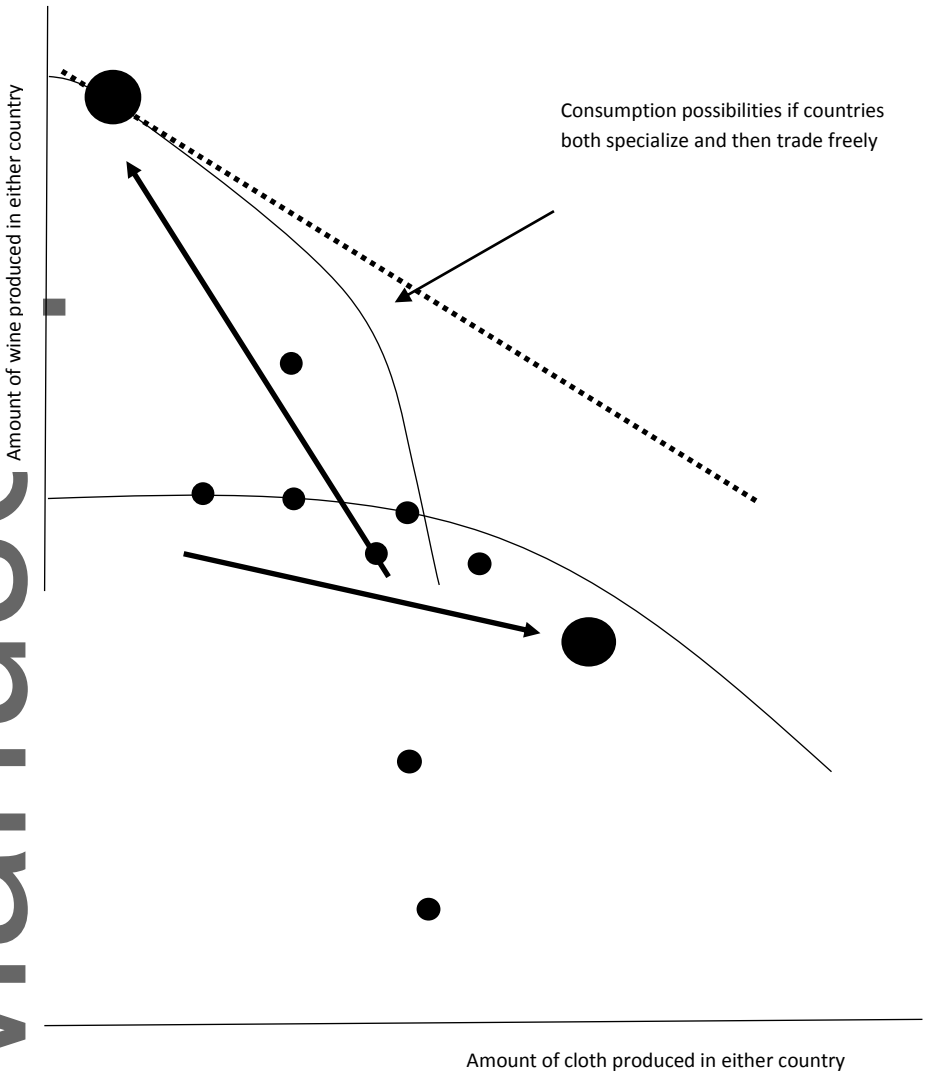

Fig. 5: Growth due to increased gains from trade over time

This article is protected by copyright. All rights reserved.

I've asked myself for a long time what might this look like in reality. It's fairly standard, even hackneyed, to ask what it looks like in a textbook. But what does it look like in the real world out there? And so the first thing I'm going to talk about that I've worked on is a setting where agriculture was very important, where it was plausible that Ricardian forces were at work, and where these regions started to trade more with one another. The particular case I'll be talking about is British colonial India during the railroad building era, in which a dense network of railroad lines was built connecting all the smaller subunits of colonial India. In this regard, I'll be talking about 235 such subunits that in British India were called districts. The period for which data exist and railroads were still unfolding is 1870-1930. The setting, therefore, is that we know of two regions – or 235 regions, really – which had this advancing opportunity, thanks to the railroads, to trade more with one another. Do they look productively more efficient throughout that time?

Figure 6 gives a sense of the improved trade opportunities in colonial India throughout this time period. Before the railroads were built, trade was confined to the coastlines and the relatively small number of navigable rivers (Panel A). By 1930, by contrast, every one of the 235 districts had a rail line (Panel B).

Author Manuscript

This article is protected by copyright. All rights reserved.

Panel A: 1853

Panel B: 1930

Fig. 6: The spread of railroads in India, 1870-1930

This article is protected by copyright. All rights reserved.

There's a lot going on here. Prices are changing, trade flows are changing, output quantities are changing, and economic welfare is changing. I've discussed a lot of that in Donaldson (2018). Here, I'd like to zoom in on the main punchline that relates to today's talk, which is shown here in regression results in column 1 (Table 2). This is a difference-in-difference regression with district and year fixed effects run across the 235 districts for the 60 years. There's a statistically precise sense in which there's a strong association between the arrival of railroads in a district and improvements in total agricultural income (measured as the value of production divided by the consumer price index over those same goods). To be more precise, there's an approximately 25% increase in agricultural income in the average district when it gets access to a railroad line. We should, of course, be skeptical about whether this implies the railroads caused this improvement, but I think there's a lot of things in the data that are consistent with this being an unbiased estimate of the true causal effect of such improvement in the railroad network.

Column 2 of the table does something slightly different. It says, well, if trade were to explain this, then surely controlling for trade would make the relationship go away. The regressor that's added in the second column does exactly that. Trade flows themselves are not observed particularly well and would be endogenous even if they were, so what I use here instead is a prediction of trade flows based on a particular Ricardian model (based on Eaton and Kortum, 2002) which accounts for things seen in other aspects of the data involving prices and trade flows and rainfall. Column 2 asks: if we try to control for the Ricardian motives for trade, does the coefficient on railroad access fall? And it does, by about half. In a very crude sense, that suggests that this Ricardian model is capable of

Author Manuscript

This article is protected by copyright. All rights reserved.

accounting for about half of the impact of railroads in the case of India. That doesn't prove that Ricardian trade (or even necessarily any other model of motives for trade) had anything to do with this. But it is certainly consistent.

This then begs the question of, given those PPFs we saw earlier, how large would you have expected the gains from trade due to railroad access to have been in a setting like this? Now, that is a hard question, because we economists don't ever have data on PPFs, just as we don't ever have data on demand curves and supply curves. What we have instead is data on the quantity produced, the quantity consumed, and the prices that prevail for a product at a given point in space and time. A similar problem arises here, but the more I thought about this in the case of the Ricardian model, the more worried I got.

Why do I say this? The logic of Ricardo's four numbers is that you gain in a trading equilibrium, relative to an autarkic one, exactly from not doing what you are bad at doing. So England is not producing wine in the Ricardian example, and Portugal might also not be producing cloth in the Ricardian example. But if an econometrician named David Ricardo came along and asked, "how efficient is England at making wine?", and he only observed the trading equilibrium, he wouldn't be able to know the answer. Indeed, it is a feature inherent to the Ricardian model that anywhere outside of autarky the econometrician could not hope to know Ricardo's four magic numbers. This is, of course, what econometricians call an identification problem.

Author Manuscript

This article is protected by copyright. All rights reserved.

Table 2: Railroads associated with increased income, India, 1870-1930

Log of real agricultural income, corrected for rainfall	(1)	(2)
Railroad in district	0.258	0.124
	(0.050)	(0.050)
"Trade share" as computed in model		-1.587
		(0.177)
Observations (number)	7 086	7 086
R^2	0.835	0.844

Source: Donaldson (2018, p930)

I don't think the identification problem in the Ricardian model is particularly unique in that regard. This is nothing more than a selection problem that pervades lots of what we as economists do, but the difficulty here is that it cuts to the heart of the Ricardian model because the gains from trade in this model come entirely from selection. So if we want to speak about how big we think the gains from trade are, we better know a lot about selection – yet this is hard when we can't usually see the differences in productivity that drive the selection process. As an aside, this problem also comes up in the labor economics literature, for example, where they call the Ricardian model the Roy model (Roy, 1951). In this case, if you were to ask yourself who will become a dentist and who will become an economist, the Roy model can predict the outcome if the econometrician using the Roy model knows how good everybody in the economy is at being an economist and at being a dentist. But (thankfully) modern labor markets are in a trading equilibrium and so (thankfully) none of us do both

Author Manuscript

This article is protected by copyright. All rights reserved.

of these tasks at the same time. And so we can't really hope to use the Roy model because we can't hope to observe the necessary primitives.

Because that key identification problem cuts right to the heart of the Ricardian model, Arnaud Costinot and I got interested in the work of agronomists, and of economists using the work of agronomists, because in some sense a major goal of agronomists is to be able to tell farmers about the farmers' PPFs, the options available to those farmers. An example of such work by agronomists comes from the Global Agro-ecological Zones Project (GAEZ) of the FAO that estimates the suitability of every small land parcel on Earth at producing different crops. Figure 7 gives a sense of this work. It zooms in on the relative productivity of two crops, wheat and cotton, in the thousands of grid cells studied by GAEZ in the case of the United States.

So we have two goods, but far more than two "countries". Each pixel in this picture is like a country, which of course would make Ricardo weep with joy because this figure is screaming at us that there is a huge scope for Ricardian gains from trade in this world,. There are regions that are really good at wheat and really bad at cotton, and there are other regions that are good at cotton and really bad at wheat. And that's just two crops, but there are lots of other crops, and hence lots of complexity to which this doesn't do anywhere near enough justice, but you can see the basic idea that we know we live in a heterogeneous world. These plots of land are all different, and these differences give rise to potentially huge gains from trade in a Ricardian world.

Author Manuscript

This article is protected by copyright. All rights reserved.

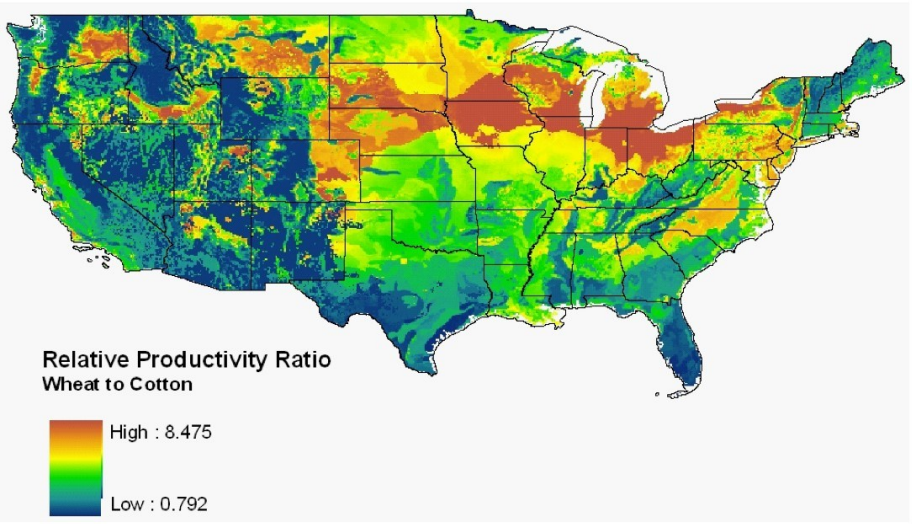

Fig. 7: Suitability estimates for wheat relative to cotton

So, I'm going to talk now about some calculations I've done with Arnaud Costinot (Costinot and Donaldson, 2016) to try to answer the question: how big are the gains from agricultural trade within the United States? But instead of quite doing that, we're do something just a little more modest, which is to ask how much those gains have been exploited so far, meaning from when the data started in 1880, until about 1997. The data consists of a panel of roughly 2600 US counties, with a median number of identified grid cells per county of 26, that can be tracked through that time. The basic intuition of what we did was very simple – we asked whether these GAEZ estimates are possibly wrong.

Author Manuscript

This article is protected by copyright. All rights reserved.

Now of course if the GAEZ predictions about the PPF are completely wrong (or completely off in some particular time period for which GAEZ didn't hope to be accurate) then this knowledge is no help. But what Arnaud and I worked out is that if the GAEZ predictions are wrong in a particularly structured way, then we can use auxiliary data to derive a set of corrections. The basic insight is that, thanks to U.S. Census data at the county level, we know both the total volume of production and the total amount of land that was required to achieve that production, for every crop and every Census year (from 1880 onwards). As long as the GAEZ model is correc about the nature of within-county-year comparative advantage (across grid cells within the county) then, with the help of the two pieces of Census information, you can adjust the PPFs inwards, outwards, and with rotations, as needed to match the Census data and the GAEZ information on comparative advantage. We have done this for 16 leading crops (plus an outside land use activity) for which data were available for the 120-year period, resulting in a 17 dimensional PPF for each county in each year.

That's step one: adjust the GAEZ estimates in order to match the output and land use under the assumption that farmers are price-taking and profit-maximizing. The second step is to acknowledge that the slope of the PPF at the point at which production is actually occurring must be the vector of relative farm gate prices that those (rational) farmers are facing. That of course is easy to figure out once you know the PPF and where production is taking place: it's the dot on the dotted curve shown earlier (Figure 4). The slope there tells you the relative price for that county and that year, if you are doing it for one county in one year.

Author Manuscript

This article is protected by copyright. All rights reserved.

Now, instead of asking a hypothetical question about what would happen if that county were able to trade freely, we ask something a little bit more mundane. Take that same county in the future (call that time period *t'*) and ask: if I were to give the county in year *t* the farm gate prices from the future *t'* how much more, or less, productive would year *t* have looked? Recall, also, that distortions are only defined in reference to relative prices, so what we care about is the farm gate prices relative to the national wholesale price, in other words a sort of wedge between the farmers' price and the free trade price. We are, therefore, going to look at the county-specific wedges and how they change over time. But we can estimate one of these wedges for every crop, county and year. And then we can ask how much more or less valuable aggregate, national production in year *t* would have been if the farmers instead faced the wedges from year *t'.*

Note that there is nothing in this calculation that necessarily implies that there will be gains from moving from year *t* to year *t'.* It could be that the wedges that rationalize the production choices are actually growing over time, opposite to the hypothetical case in Figure 5 where the production bundles in each country are moving towards the more efficient equilibrium.

Author Manuscript

This article is protected by copyright. All rights reserved.

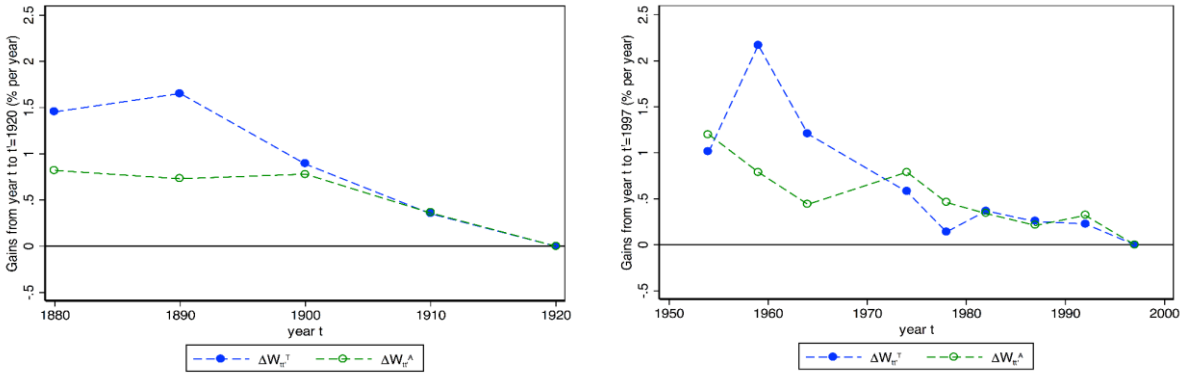

a) Early period, 1880-1920 b) Late period, 1954-1997

Note: Upper line with solid dots: gains from between-county integration (as if a trading technology). Lower line with hollow dots: gains from within-county PPF movements

Fig. 8: Estimated gains from internal integration of US agricultural markets

What we do find is illustrated in Figure 8. Panel (a) shows our estimates for 1880 through 1920, and panel (b) the case covering 1954 to 1997. There are some data concordance problems that prevent us from connecting 1920 to 1954 without further assumptions, so we have not done that. However, the upper line (with the solid dots) shows at each point an estimate of the appropriately weighted US-wide answer to my previous question: if you were to take that county out of year *t* and into some year *t'* in the future, how big would the gains be? In all these pictures, the last year is always going to

Author Manuscript

This article is protected by copyright. All rights reserved.

be the year *t'*. Thus, in the example of panel (b) year *t'* is 1920. So obviously if you take 1920 and you replace 1920's price wedges with itself, you should get no gains, which is precisely why the solid dot on 1920 is 0. By contrast, the solid dot for 1880 is definitely not zero, but positive, suggesting that there are gains from this hypothetical movement in wedges. In fact, the gains are large, namely a 1.5 percent compound annual growth rate over 40 years, i.e. 1.5% of real productivity growth per year.

Figure 8 also shows that this rate of aggregate productivity growth is typically falling over time, suggesting that the wedges are getting smaller over time. Nevertheless, they are always big in some meaningful sense: 1.5% real productivity growth per year is nothing to sniff at. Another way of putting that in context is to ask something different, shown in the lower line with the hollow dots: in this case, we don't change wedges; instead, we hold prices constant and replace the county's year-*t* PPF with its year-*t'* PPF. So it's exactly opposite as the previous case. The upper line holds the PPFs fixed and asks how much more allocative efficiency has there been thanks to price convergence or wedge convergence, while the lower line holds the prices fixed and asks by how much must productivity grow. This is probably an exercise you are more familiar with when you think about productivity-- PPFs moving out is undeniably productivity growth. The lower line shows you that number and again, it's big, about 1% per year. But interestingly it's of the same rough magnitude as the gains from enhanced integration.

Author Manuscript

This article is protected by copyright. All rights reserved.

One caveat to these results is that calculating the cost of these distortions, as we have done here, and giving you an estimate, depends on how you want to view such price wedges. The upper line represents a scenario in which not only are prices being distorted such that two farmers are not facing the same relative price, but real resources are also being used up (i.e. productive factors are employed) in the act of getting goods from one location to another.

This is not a crazy metaphor, because transport, of course, uses real resources. However, from the agricultural sector point of view, it is not clear that we would want to consider gains from transport cost reductions as gains within the agricultural sector. Therefore, in Figure 9 the line with the solid dots repeats the calculation for a scenario in which there is no real resource loss at all from the wedges themselves. This would be like the hypothetical setting where these wedges correspond purely to *ad valorem* taxes that generate revenue that is redistributed in a lump sum fashion without any further loss.

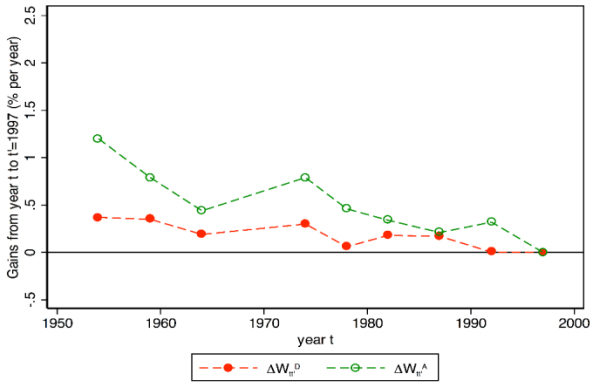

This article is protected by copyright. All rights reserved.

a) Early period, 1880-1920 b) Late period, 1954-1997

Note: Line with solid dots: gains from between-county integration (as if a tax). Line with hollow dots: gains from within-county PPF movements

Fig. 9: Estimated gains from internal integration of US agricultural markets

Consider, then, the latter scenario about these wedges – they just distort prices, and do nothing else. These gains from integration are certainly lower, but still large and broadly in the same ballpark as the upper line (with the hollow dot), which, as before, is the pure productivity gains.

This result surprised us, but then we looked again at the map in Figure 7. No Ricardian thinker could look at this map and think that the gains from trade are not real. The only question is when did the trading start? When did the gains start to be enjoyed? Our estimates imply that even by 1880 the gains were not fully exhausted. There was still lots on the table, and over the ensuing 120 years a lot of US agricultural growth can be attributed purely to homogenization of prices.

Author Manuscript

This article is protected by copyright. All rights reserved.

2. Trade and volatility

The second area I want to talk about is volatility. Let me start by apologizing for how I will use that term. There is not going to be here a fully consistent view of randomness and uncertainty and the kind of behavior that people engage in such uncertain environments. I have skipped many of those complexities, but I will show you what I can say.

To start, let's go back to Figure 4, but now imagine we have two countries that have exactly the same PPF (see Figure 10a). They are both producing and consuming at point A because there is no need to trade – there are no gains from trading. Imagine now that these two economies are subject to some random shocks to their productivities, maybe due to the weather that produces the ingredients of the wine or the ingredients of the cloth. To make the example easier imagine that one country gets a good shock on one of its goods and a bad shock on the other. That means the PPF will not move strictly inwards when you get a bad shock or strictly outwards when you get a good shock – it will just tilt. Now imagine that it just so happens that the two countries' PPFs tilt in opposite ways. Whenever one gets a good shock, the other gets a bad shock (Figure 10b).

In any state of the world, whether it is the no-shock state or the I-get-the-left-shock-and-you-get-the-right-shock state or vice versa, we're still able to trade (Figure 10c), in other words the ability to trade can act like insurance. Obviously without a literal model of risk aversion and a notion of state

Author Manuscript

This article is protected by copyright. All rights reserved.

contingent assets, etc., - insurance markets if you like - it's a little hard to talk about this formally, but you get the intuitive idea that one of the advantages of trade might just be that we don't use it often because we don't need it. In the average 'state' of the world, no one wants to trade. Yet there are times when you get a bad shock and you want to trade, or your neighbor gets a bad shock and she wants to trade, or you both get shocked (in an uncorrelated way) and you want to trade even more, and that's all that's going on in this example.

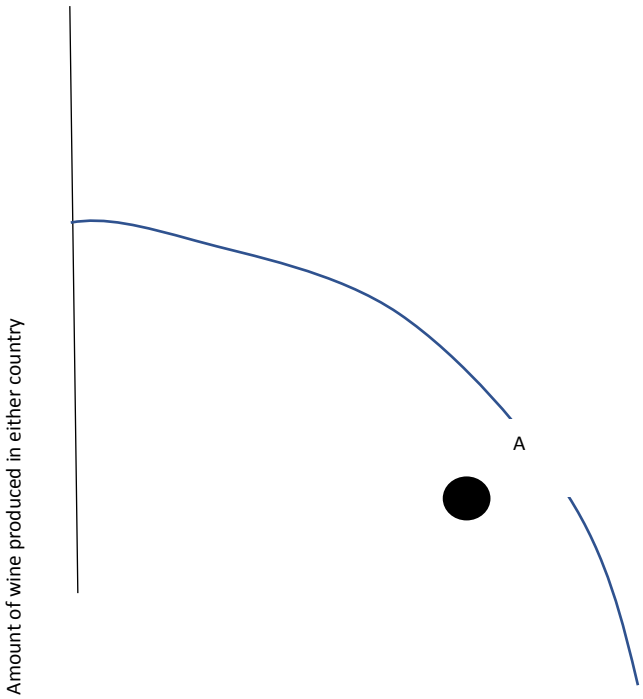

Author Manuscript

This article is protected by copyright. All rights reserved.

Author Manuscript

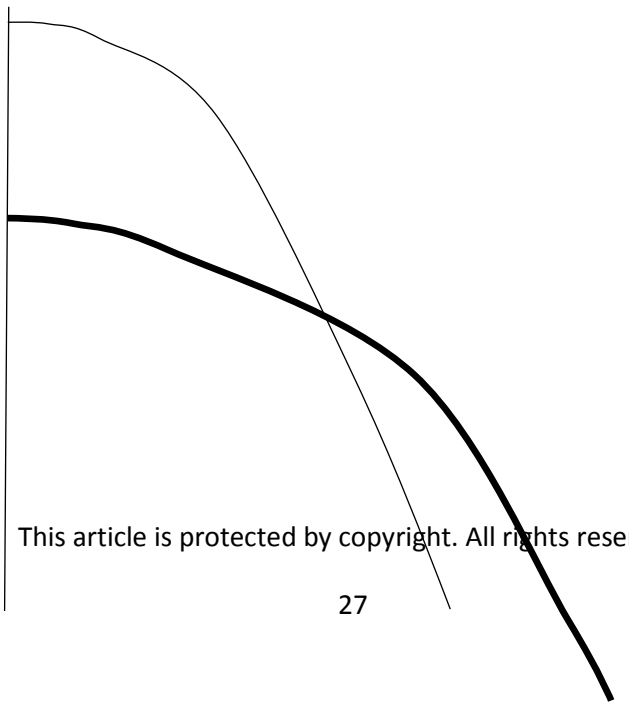

Figure 10a: Trade and volatility

This article is protected by copyright. All rights reserved.

Author Manuscript

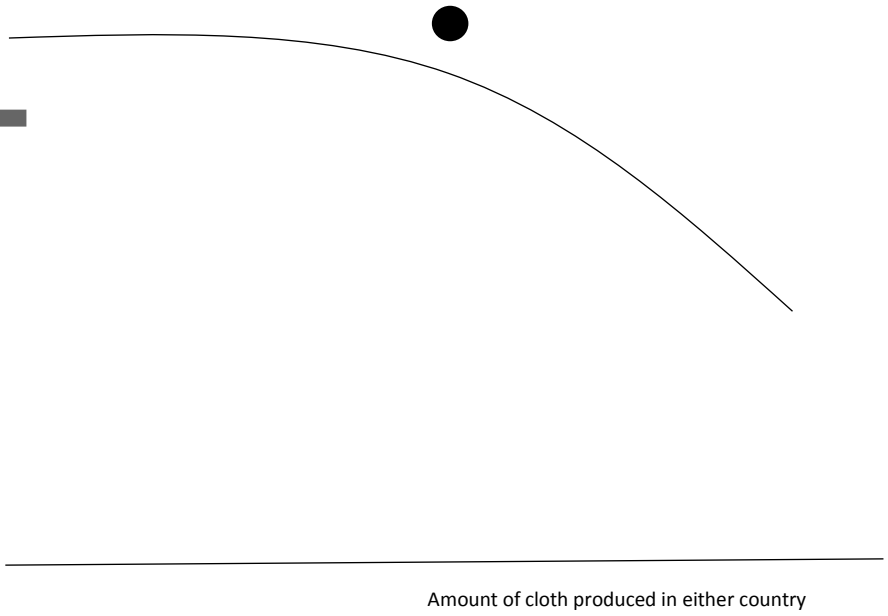

Fig. 10b: Trade and volatility

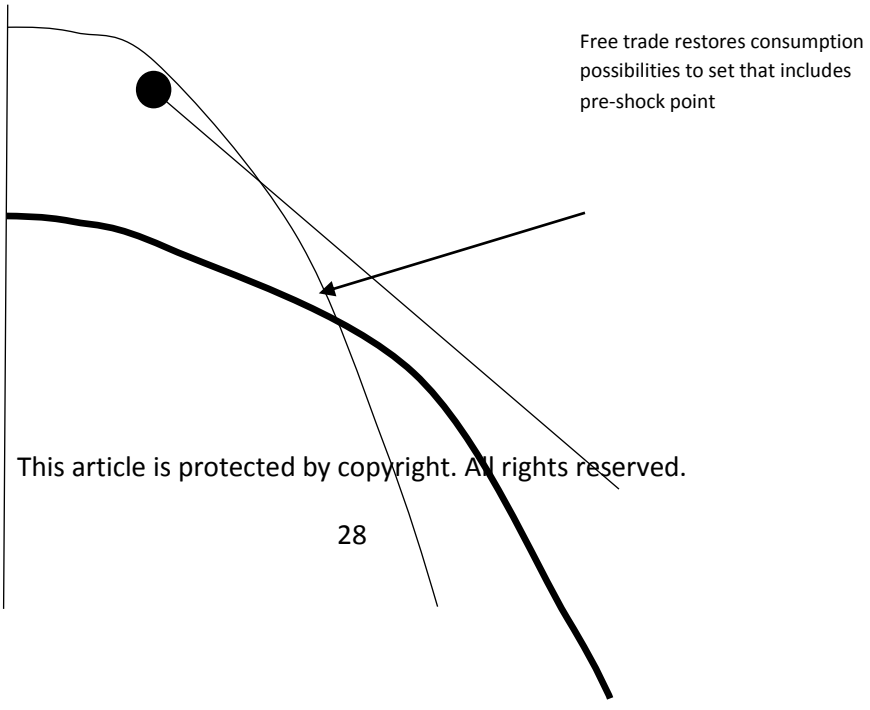

This article is protected by copyright. All rights reserved.

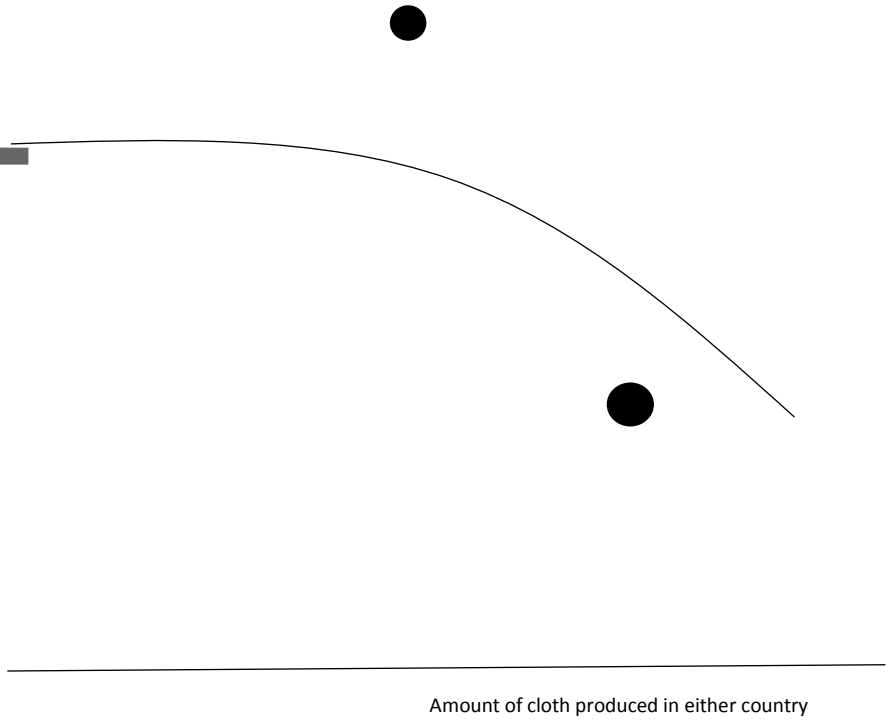

Fig. 10c: Trade and volatility

What does this look like in reality? I'll start by returning to colonial India. This setting, as Will alluded to, is no stranger to agricultural volatility and its consequences. The historian Mike Davis called this the era of the late Victorian holocausts because of the famines that occurred, and one view of those famines that they had at their root in a set of productivity shocks that were the result of random events like the weather (Davis, 2000). This is important because I'm going to focus on one bit that's

Author Manuscript

This article is protected by copyright. All rights reserved.

random, that due to rainfall variation, which was such a big deal that in the 19th century the British were collecting daily rainfall data at an incredible number of weather stations – fully 3600 of them in fact.

Figure 11 shows the rainfall over the 40 years between 1880 and 1920 in six major provinces of colonial India. At the top left is Bengal, where rainfall is substantial (around 1.5 meters in most years), yet with some clear volatility – volatility that would, of course, look even greater if we dug down to each of the districts of this province. And if you look at the next province, Bombay, there seems to be again a great deal of volatility, and this is volatility that is definitely not perfectly correlated with that in any of the other provinces shown. Therefore, the different regions are getting rainfall shocks, and these are not perfectly correlated. This suggests the possibility that trade could help, and that is what Robin Burgess and I looked at. The results are given in Table 3 (Burgess and Donaldson, 2010). In this study, we drew on the work of a wonderful book about the famines of the era by Srivastava (1968) who coded up every district of the country in every year with a qualitative index between zero and three, giving four different values of famine severity, with zero as no famine. We used this ordinal index as the dependent variable in an ordered logit model, and studied some of its causes.

The first column shows that, on average, some of these famines had at their root random events such as lack of rainfall. There is, in fact, a strong negative correlation between rainfall and famine

Author Manuscript

This article is protected by copyright. All rights reserved.

severity. Therefore, while there may be many other causes of famine, rainfall was at least one of them on average. The next column addresses whether railroads are correlated with famine getting less severe. Remember the railroads are being built throughout time, so this is always a difference-in-difference regression: in a given district, are you less likely to see a famine after you get access to railroads? The positive sign on the coefficient tells you that when the average railroad entered the average district, the same rainfall shock was less likely to lead to famine. That supports Will's contention: this is consistent with the idea that the railroads were mitigating famine, subject to the usual concerns about attributing causation to correlation.

The likelihood of seeing a famine for a given amount of rainfall after you have access to railroads is calculated by adding the two numbers here (-2.2 and positive 1.8), which of course gives you a number that's almost zero. This tells you that on average, a district that had access to the railroads was less likely to be subject to famines given the same rainfall shocks. Of course, if there are other causes of famine, such as may have been the case in Bengal in 1942-43 as Amartya Sen has argued (Sen, 1981), then this regression would not pick up those causes.

Author Manuscript

This article is protected by copyright. All rights reserved.

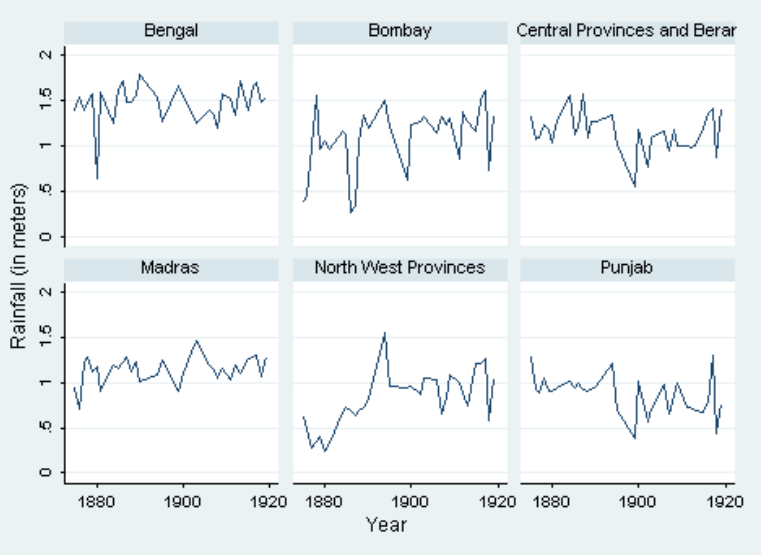

Source: Burgess and Donaldson (2017: 38)

Fig. 11: Rainfall volatility in India, 1880-1920

Table 3: Ordered logit model using the "famine severity index" (Srivastava, 1968)

Dependent variable: famine severity index	(1)	(2)	(3)	(4)	(5)
Railroad in district (0,1)	0.194	-1.625***	0.309	-2.178***	-2.136***
	(0.374)	(0.572)	(0.390)	(0.690)	(0.754)
Rainfall in district, year *t* (meters)	-0.855***	-2.218***	-0.860***	-2.316***	-17.35
	(0.208)	(0.532)	(0.204)	(0.518)	(20.40)
(Railroad in district) x (rainfall in district, year *t*)		1.858***		1.848***	1.729***
		(0.541)		(0.521)	(0.565)

Author Manuscript

This article is protected by copyright. All rights reserved.

Rainfall in district, year *t*-1 (meters)			-0.699***	-1.171***	9.316
			(0.215)	(0.395)	(21.51)
(Railroad in district) x (rainfall in district, year *t*-1)				0.692*	0.758*
				(0.404)	(0.458)
(Rainfall in district) x (trend)	No	No	No	No	Yes
Observations	3809	3809	3551	3551	3551
Pseudo R^2	0.248	0.260	0.255	0.271	0.271

The final thing I will talk about in this theme of "volatility" is climate change. This is probably not the right way to think about the broad patterns of climate change, but I wanted to map back to the way that we have so far today been thinking about comparative advantage and agricultural trade.

Start with an economy with a given PPF. Suppose it is open to trade. But now shock the PPF in a certain way, as we might expect to occur as a result of climate change. The last question I'll cover today is: how does the ability to trade affect the extent to which a region will be affected by this climate change-driven change in the PPF? The predicted percentage change in any crop can be calculated in an agronomic model (such as the GAEZ model) if we simply feed in the climate change

Author Manuscript

This article is protected by copyright. All rights reserved.

predictions from a climatological model (such as one from the Hadley Centre[2]). Figure 12 illustrates these predictions for relative wheat and rice yields around the world. It is clear that these relative yields are predicted to move in heterogeneous ways. So the climatologists and agronomists expect a vast and complicated set of changes to Ricardian comparative advantage over the next century.

In work with Arnaud Costinot and Cory Smith (Costinot, Donaldson and Smith, 2016), we tried to ask the very simple question: if you took those numbers as well as similar numbers for every other major crop and fed them into a standard trade model, where the spatial units of analysis are the grid cells of Figure 12, how big would the effect of that tilting of the PPF be?

To address this, we used a standard and simplified trade model with three key elasticities. The first of these is the Armington differentiation elasticity (σ), a source of cross-country differentiation in production. Of course, consumers don't want to consume a fixed relative amount of rice and wheat at any price; instead, they might want to substitute between crops and so we call that elasticity of substitution κ. Finally, we enriched earlier work (such as that in Costinot and Donaldson, 2016) by allowing for within-grid cell heterogeneity in land quality for each crop. This amounts to having a PPF for each grid cell that is curved, rather than flat. We let this curvature – this cross-crop elasticity of

[2] https://www.metoffice.gov.uk/weather/learn-about/climate-and-climate-change/climate-change/hadleycentre

Author Manuscript

This article is protected by copyright. All rights reserved.

transformation, if you like – within grid cells be denoted θ. Grid cells still difer in terms of the location and tilt of their PPFs, in line with the GAEZ model, but now there is additional within-grid cell heterogeneity.

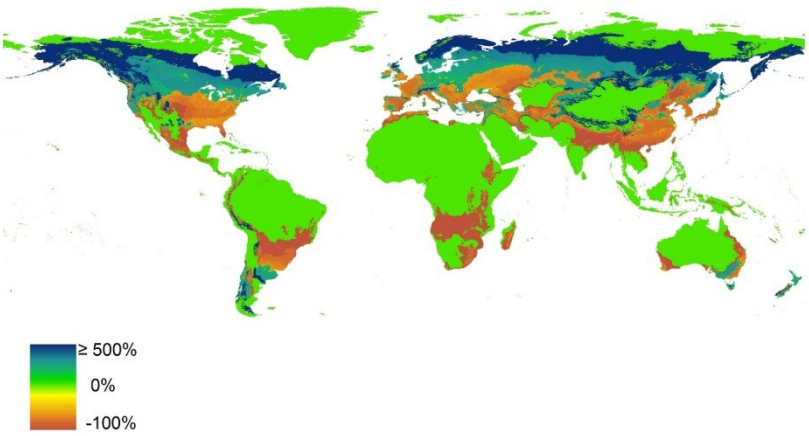

Fig. 12: The predicted change in relative wheat-to-rice yields according to the GAEZ model under a climate change scenario (Hadley CM3 A1F1)

Author Manuscript

This article is protected by copyright. All rights reserved.

So we estimated those three parameters (σ, κ, and θ) using cross-country and cross-crop data. Estimating these demand- and supply-side elasticities of course requires instrumental variables to circumvent simultaneity bias. On the demand side we derived this from the GAEZ model, which makes predictions about how the exogenous geographical characteristics of a location map into the location's productivity and hence (inversely) its prices (as long as trade costs are not zero). This suggests a host of instruments one can use, and we used the simplest ones we could think of. The FAOSTAT[3] trade, production and price data allow us to do all of this, and the elasticities are important for the estimates that we would compute in general equilibrium.

These elasticities are reflected in Table 4. The Armington elasticity is around 5, which suggests that there are barriers to within-crop substitution across countries. The κ elasticity of substitution between crops is more in the ballpark of 3, and the θ parameter that governs the production substitution within grid cells across crop activities is low, at around 2.5. This is certainly not a flat PPF.

[3] See http://www.fao.org/faostat/en/#home

Author Manuscript

This article is protected by copyright. All rights reserved.

Table 4: Demand parameters estimated via IV, using supply-side GAEZ variation

Parameter	Estimate	Method	Observations	First-stage F-statistic
σ	5.40 (1.26)	IV	4120	-0.06 (19.64)
κ	2.82 (0.54)	IV	408	-0.12 (22.3)
θ	2.46 (2.28, 2.62)	NLS	379	NA

Notes: In the first row (estimates of σ) the standard error (SE) and first-stage F-statistic are clustered at the crop-exporter and crop-importer levels. In the second row (estimates of κ) the SE and first-stage F-statistic are clustered at the importer level. And in the third row (estimates of θ) the 95% confidence interval is estimated via a bootstrap procedure with replacement at the country level.

We then go on to use these estimated elasticities to quantify some of the general equilibrium effects of climate change in the agricultural context. One way of seeing these estimates is shown in Figure 13. I must emphasize that if you were really trying to predict the consequences of climate change for these countries in these kinds of goods, then there are tons of bells and whistles you would probably want to add to the model we are using here. A lot of good work does that, and this work has definitely cut some corners with respect to that goal. What we were interested in here is to ask: in this standard trade model, holding the standard trade model features constant but adding the spatial granularity coming from GAEZ (as seen in Figure 12), how big would you expect the effects of climate change effects to be?

Author Manuscript

This article is protected by copyright. All rights reserved.

So here are those estimates. Start with panel (a) of Figure 13. Each dot in this picture is a country.

The numbers on the x-axis are the ones you'd probably want to think about first. Here full adjustment is allowed, and the results are probably not shocking to you. They show that, because of climate change, an average country will experience something between 0 and 15 percent loss in the economic welfare it derives from the agricultural sector. In some countries, it's closer to 0, while in some countries, it's very bad. But typically the damage is in the range you might have expected. The numbers on the y-axis of this figure are the gains in a no adjustment scenario. This is where you force farmers to keep growing what they're growing. Of course, they would want to change what they do when there are movements in their PPFs, the prices they receive, or both. But just as a thought experiment you could hold those things constant. Of course the losses under this scenario are bigger for most countries, as we'd expect them to be when farmers can't adjust to the shock to their surroundings.

Panel (b), on the other hand, shows results that are perhaps more surprising. Here we asked what will happen if we allow the farmers to adjust, but not the traders. By that I mean that a country's export share of any crop cannot increase. In this case the scatter plot is right on the 45 degree line, implying that for a typical country, climate change is going to hurt, but if you forbid traders to respond, that is if a country's trade patterns have to remain the same as they were before the shock, then this additional constraint is almost irrelevant. This finding surprised my coauthors and me. But maybe it shouldn't have because Figure 12 should make it clear that a lot of the shock and, hence, a

Author Manuscript

This article is protected by copyright. All rights reserved.

lot of the adjustment that needs to occur, can take place within countries. This resonates with themes we covered earlier. If you think about the Ricardian model at this sub-national level, countries start to become very heterogeneous objects and gains from trade and opportunities for adjustment within countries are important because of the heterogeneity that resides within countries.

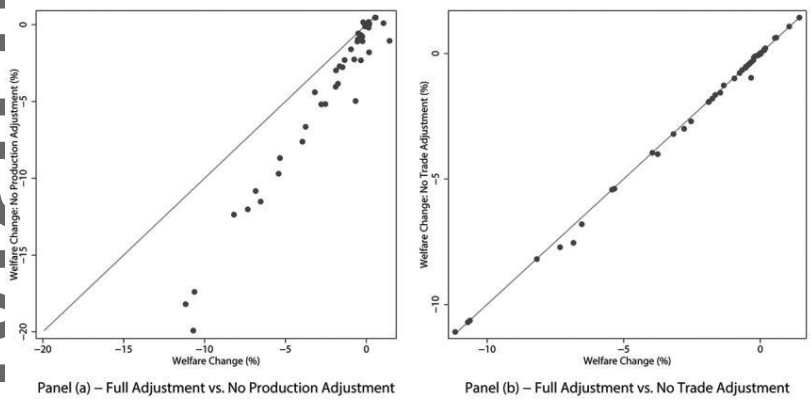

Fig. 13: Comparison of simulated welfare changes under alternative assumptions about adjustment

Author Manuscript

This article is protected by copyright. All rights reserved.

3. Concluding remarks

I'm going to offer just a few remarks and then let you move on to the many more interesting things that this conference has to offer. I for one have been in awe of Ricardian theory since I first encountered it. I've found that those four magic numbers go a long way. They can become millions of magic numbers in the kinds of settings I've shown you here. They can become an infinite number of magic numbers if we allow for smooth geographies. For these reasons and more they've been appreciated by economists for centuries now. Yet as great as this appreciation is, I think it's fair to say that the quantification of those forces has lagged behind. And that's not a surprise; quantification in general is hard, but it might be especially hard in this Ricardian model because of the identification problem I mentioned that's at the very heart of the model.

All I've tried to do today is show you a few examples of how, in an agricultural context, the Ricardian model is particularly easy to work with, relatively speaking, thanks to the work of agronomists. And then how this model can be applied to examine things like the gains from trading, in both the static sense and the volatility sense, because of the amazing heterogeneity that is out there in the world around us, and that we see revealed in agricultural settings. That suggests that international heterogeneity as well as intranational heterogeneity can be big. That, in turn, suggests that the gains from domestic market integration can be big, and that the domestic forms of adjustment to shocks – be they of high frequency like food supply shocks, or low frequency like climate change – should not be underestimated.

Author Manuscript

This article is protected by copyright. All rights reserved.

Bibliography

Burgess, R., & Donaldson, D. (2010). Can openness mitigate the effects of weather fluctuations? Evidence from India's famine era. *American Economic Review: Papers and Proceedings, 100,* 449-453. http://aeaweb.org/articles.php?doi=10.1257/aer.100.2.449

Burgess, R., & Donaldson, D. (2017) Railroads and the Demise of Famine in Colonial India. Working Paper. Available at: http://www.lse.ac.uk/economics/Assets/Documents/personal-pages/robin-burgess/railroads-and-the-demise-of-famine.pdf. Accessed 5 September 2019

Costinot, A., & Donaldson, D. (2016). How large are the gains from economic integration? Theory and evidence from US agriculture, 1880-1997. NBER Working Paper 22946, National Bureau of Economic Research, Cambridge MA.

Costinot, A., Donaldson, D., & Smith, C. (2016). Evolving comparative advantage and the impact of climate change in agricultural markets: Evidence from 1.7 million fields around the world. *Journal of Political Economy, 124,* 205-248. https://doi.org/10.1086/684719

Davis, Mike, 2000. *Late Victorian holocausts: El Niño famines and the making of the Third World*. UK: Verso

Donaldson, D., (2018). Railroads of the Raj: Estimating the impact of transportation infrastructure. *American Economic Review 108,* 899-934. https://doi.org/10.1257/aer.20101199

Eaton, J., & Kortum, S. (2002). Technology, Geography and Trade. *Econometrica*, 70, 5, 1741-1779. https://doi.org/10.1111/1468-0262.00352

Author Manuscript

This article is protected by copyright. All rights reserved.

Eaton, J., & Kortum, S. (2012). Putting Ricardo to work. *Journal of Economic Perspectives, 26,* 65-90. http://dx.doi.org/10.1257/jep.26.2.65

Roy, A. D., 1951. Some thoughts on the distribution of earnings. *Oxford Economic Papers (New Series),* 3, 135-146. https://doi.org/10.1093/oxfordjournals.oep.a041827

Samuelson, Paul A. (1972), "The Way of an Economist." Reprinted in The Collected Papers of Paul A. Samuelson. Ed. R. C. Merton. Cambridge: Cambridge MIT Press.

Sen, Amartya, 1981. *Poverty and famines: an essay on entitlement and deprivation*. Oxford: Oxford University Press

Author Manuscript

This article is protected by copyright. All rights reserved.