Agricultural Productivity, Trade Costs, and Structural Change*

Martin Fiszbein§
Boston University and NBER

Will Johnson¶
Dartmouth College

January 2020

VERY PRELIMINARY AND INCOMPLETE

click here for latest version

Abstract

Agricultural productivity growth has two opposing effects on structural change. On the one hand, it increases income; given non-homothetic preferences, this pushes labor into non-agriculture. On the other hand, it shifts comparative advantage toward agriculture. The relative strength of these two forces depends on trade openness. We provide reduced-form evidence consistent with heterogeneous effects of agricultural productivity depending on levels of openness. Exploiting the Green Revolution as a source of plausibly exogenous variation in agricultural productivity, we find that for 18%-37% of the countries in our sample, agricultural productivity has negative effects on structural change. We develop an Eaton-Kortum model with non-homothetic preferences that features heterogeneous effects of agricultural productivity by levels of openness. We calibrate the model and show that it can explain a sizable portion of the observed changes in countries' sectoral labor shares from during the Green Revolution 1961 to 1995.

*We thank Treb Allen, Diego Comin, Ryan Edwards, Jim Feyrer, Stefania Garetto, Dilip Mookherjee, and Nina Pavcnik for useful comments and suggestions. Arthur Smith provided excellent research assistance. All errors are our own.

§fiszbein@bu.edu.

¶william.Johnson-2@dartmouth.edu

1 Introduction

Agricultural productivity growth has two opposing effects on structural change. On the one hand, it releases labor from the agrictural sector, inducing structural change. On the other hand, it shifts comparative advantage toward agriculture. The relative strength of these two forces depends on trade costs.

Both classic and recent contributions emphasize *positive* effects of agricultural productivity on industrialization (the *labor push* effect). Nurkse (1953) and Rostow (1960) emphasized that significant increases in agricultural productivity were a pre-condition for industrialization and the take-off into modern economic growth. Schultz (1953) emphasized that many developing countries' experiences were so close to subsistence that large fractions of their populations inevitably had to be employed in agricultural production—"the food problem." A number of recent contributions also argue that improvements in agricultural productivity are a key driver of industrialization (see Gollin et al., 2002, 2007; Ngai and Pissarides, 2007; Alvarez-Cuadrado and Poschke, 2011; Herrendorf et al., 2014).

However, **?** shows that for a small open economy, agricultural productivity has *negative* effects on industrialization because of the comparative advantage effect. He thus argued, "the openness of an economy should be an important factor when planning development strategy and predicting growth performance." While Matsuyama's paper became a well-known classic in macro-development, both development policies and research often assume that agricultural productivity can only have positive effects. The small open economy with free trade case may be an extreme case that is not directly applicable to the real world, but Matsuyama's insight may still be relevant.

Intuitively, real world economies are intermediate cases lying between the polar extremes of Matsuyama's model—free trade and autarky—and in such cases both the labor push effect and the comparative advantage effect are relevant. Moreover, the relative magnitude of these depends on the size of transport costs, implying a source of heterogeneity in the effects of agricultural productivity on structural change that has not received much attention.

We start by documenting heterogeneity in the effects of agricultural productivity on structural change depending on trade opennes. We provide reduced-form evidence consistent with such heterogeneous effects, exploiting the Green Revolution as a source of plausibly exogenous variation in agricultural productivity. We find that for 18%-37% of the countries in our sample, agricultural productivity has negative effects on structural change.

We then show that an Eaton-Kortum model with non-homothetic preferences can capture the heterogeneous effects of agricultural productivity depending on levels of openness. Our model is a multi-country, two-sector (agriculture and non-agriculture) Ricardian model of trade. The structure of the model follows Eaton and Kortum (2002) ("EK"), but with non-homothetic CES preferences over two sectors, following Comin et al. (2017) ("CLM").[6]

This model structure allows us to go beyond the two extreme cases analyzed by **?** (a closed economy and a small open economy). We start with a simple two-country numerical simulation of the model to show that Matsuyama's point holds under cases of intermediate levels of openness as well, with the relative strength of the labor push effect and the comparative advantage effect varying in a continuous way with openness.

We then take the model to data, exploiting the Green Revolution as a natural experiment, in order to quantify the relative strength of the labor push effect and comparative advantage effect. During the Green Revolution, when many different countries were simultaneously experiencing increases in agricultural productivity − with the size of this increase varying from country to country − and these countries were trading with each other subject to a complex network of trade costs varying at the country-pair level, our structural model makes a rich set of predictions regarding how the Green Revolution

[6]Note that CLM's model is a closed-economy model; we embed it in the EK model of trade. Our model is also similar to Uy et al. (2013) ("UYZ"), except that preferences in UYZ are such that the extent of non-homotheticity varies substantially by income level (in particular, as income goes to infinity, preferences are asymptotically homothetic). CLM find evidence in both international data and sub-national data that the extent of non-homotheticity is roughly constant across income levels, which their preference structure rationalizes. Given that our analysis looks at a much broader range of countries than UYZ, this is important for our purposes.

changed the global trading equilibrium, in particular how much each country's sectoral labor shares changed.

We evaluate these predictions of the model based on a calibration exercise. We use 1961 data to calibrate the model, then feed into the model the plausibly exogenous changes in agricultural productivity that occurred from 1961 to 1995 during the Green Revolution, and then compute the new equilibrium of the model. We show that the model is able to match much of the observed changes from 1961 to 1995 in countries' sectoral labor shares. We take this as evidence that the Green Revolution was an important driver of structural change during this period, and that our structural model is successful in capturing how increases in agricultural productivity affected different countries differently depending on their position in the global trading network.

2 Reduced-form evidence

2.1 Data on non-agricultural shares, openness and productivity

Our objective is to estimate how the effect of higher agricultural productivity on structural change varies depending on a country's openness to trade. We do this in this section by running OLS and IV regressions of the following general form:

$$Y_{it} = \beta_0 + \beta_1 O_{it} + \beta_2 A_{it} + \beta_3 O_{it} A_{it} + \epsilon_{it} \tag{1}$$

where Y_{it} is a measure of the non-agricultural share of the economy in country i at date t, O_{it} is a measure of openness to trade, and A_{it} is a measure of agricultural productivity. As discussed in the introduction, and as will be made precise in section 3, theory predicts that $\beta_1 > 0$, $\beta_2 > 0$, and $\beta_3 < 0$.

We consider two measures of the non-agricultural share: the non-agricultural share of GDP and the non-agricultural share of employment. For non-agricultural shares of GDP, we use sectoral shares of value added as measured by the World Bank's World Development Indicators (WDI). For non-agricultural shares of employment, we use

data from Wingender (2014).

Both of these measures of the non-agricultural share are informative about the impacts of agricultural productivity and openness on structural change. However, one must take care in interpreting the coefficients when estimating regressions of the form (1), depending on whether shares of GDP or employment are being used. When using shares of GDP, one must bear in mind that there is a certain mechanical effect of agricultural productivity: if agricultural productivity increases, even if no other variables in the economy (including prices) endogenously respond at all, the non-agricultural share of GDP will decrease mechanically.

When using shares of employment, the effect of agricultural productivity growth on the non-agricultural labor share might depend on not only openness, but also on whether the new agricultural technology is more or less labor-intensive than the old technology. This is not captured in our one-factor model, hence this is an important caveat when mapping between the empirics and our theory.

We measure openness as the trade share of GDP, using data from the WDI. In Appendix A, we also consider export and import shares separately.

Our primary measure of agricultural productivity, following Moscona (2018), is caloric yields per acre. We use FAO GAEZ estimates of each country's production of each crop, combined with USDA tables that estimate kilocalories per kilogram for each crop, in order to compute the implied caloric yield per acre. That is, total productivity in calories per acre is given by:

$$ProdCal_{i,t} = \sum_{K} \frac{c_{i,k,t}}{a_{i,k,t}} \left(\frac{a_{k,i,t}}{\sum_K a_{k,i,t}} \right), \tag{2}$$

where $c_{i,k,t}$ is the caloric yield of crop k in country i in year t and $a_{k,i,t}$ is the acreage of crop k in country i at time t. We use data on the caloric productivity of rice, wheat, maize, sorghum, potato, groundnut, dry beans and millet.

2.2 OLS results

For our main OLS estimates, we regress the non-agricultural share of GDP (or the labor force) on the trade share, caloric productivity and their interaction:

$$Y_{i,t} = \beta_0 + \beta_1 \frac{Trade_{i,t}}{GDP_{i,t}} + \beta_2 ProdCal_{i,t} + \beta_3 \left(\frac{Trade_{i,t}}{GDP_{i,t}} \times ProdCal_{i,t} \right) + \delta_t + \gamma_i + \epsilon_{i,t}, \quad (3)$$

where δ_t and γ_i are year and country fixed effects, respectively. As aforementioned, theory predicts a positive coefficient on β_1 and β_2 and a negative coefficient on the interaction term, β_3.

Columns 1-3 of Table 1 show the results for the non-agricultural share of GDP, while Columns 4-6 show the results for the non-agricultural share of the labor force.

	(1) Non-Ag %GDP	(2) Non-Ag %GDP	(3) Non-Ag %GDP	(4) Non-Ag %LF	(5) Non-Ag %LF	(6) Non-Ag %LF
Trade/GDP	0.216***	0.189***	0.059**	0.369***	0.362***	0.023
	(0.045)	(0.044)	(0.024)	(0.060)	(0.063)	(0.020)
ProdCal	0.577***	0.509***	0.093	1.237***	1.219***	0.045
	(0.084)	(0.085)	(0.060)	(0.144)	(0.153)	(0.076)
Trade/GDP X ProdCal	-0.003***	-0.003***	-0.001**	-0.006***	-0.005***	-0.001
	(0.001)	(0.001)	(0.000)	(0.001)	(0.001)	(0.001)
Year FE	No	Yes	Yes	No	Yes	Yes
Country FE	No	No	Yes	No	No	Yes
Cluster	Country	Country	Country	Country	Country	Country
N Obs	6339	6339	6339	5817	5817	5817

Table 1: OLS Outcome measures: percent non-agricultural GDP and percent non-agricultural labor force. Openness measure: trade share. Productivity measure: caloric productivity. Panel Regression, 1961-2016.

This table shows a strong positive relationship between both trade and productivity with the non-agricultural share of GDP, and suggestive evidence of a negative relationship between the interaction term and the non-agricultural share of GDP. The evidence is less conclusive in the relationship between these terms and the non-agricultural share of the labor force, particularly when controlling for country fixed effects, perhaps due to the caveat discussed above regarding labor intensities varying across different agricultural technologies.

2.3 IV results

There are, of course, potential endogeneity issues regarding the effects of openness and agricultural productivity on sectoral GDP and labor shares. To address this, in this section we instrument for openness using the distance-weighted GDP of trading partners, and we instrument for agricultural productivity using the Green Revolution as a natural experiment.

Following Helliwell (2000), our IV for a country's openness is the distance-weighted average GDP of every country besides that country. Here we are implicitly using the logic of a gravity equation to identify expected changes in trade via changes in the GDP of a country's trading partners. We measure distance-weighted partner GDP ($dGDP$) as follows:

$$dGDP_{i,t} = \sum_{j \neq i} \left(GDP_{j,t} / dist_{i,j} \right), \tag{4}$$

where $GDP_{j,t}$ is the GDP of country j in year t and $dist_{i,j}$ is the physical distance between countries i and j, using data from Feyrer (2019). This measure satisfies the relevance requirement given the logic of gravity equations; that is, the higher the GDP of nearby countries (or alternatively, the closer high-GDP countries are), the higher a country's trade share will be. The exclusion restriction is satisfied under the assumption that changes in partner GDP are exogenous. That is, the identification assumption is that changes in other countries' GDP affect a country's sectoral shares only through trade with that country. Although this assumption almost certainly does not hold perfectly, the distance-weighted average GDP of other countries is likely much less endogenous to a country's sectoral shares than within-country drivers of a country's openness, e.g. the quality of a country's institutions.

For productivity, we use the Green Revolution as the source of our IV. In particular, we follow Moscona (2018) by measuring the predicted productivity, $P_{i,t}$ as the average (across crops) of the maximum agro-climatic productivity of crop k at time t in country i, based on the availability of high-yield variety (HYV) crops of type k at time t

combined with the climatic conditions of country i:

$$P_{i,t} = \frac{\sum_k \left((1 - I_{k,t}) P_{i,k}^L + I_{k,t} P_{i,k}^H \right)}{N_{i,t}}, \tag{5}$$

where $I_{k,t}$ is an indicator of the presence of HYVs for crop k at time t, $P_{i,k}^L$ and $P_{i,k}^H$ are the potential yields for crop k in country i under the low- and high-input measures respectively and $N_{i,t}$ is the total number of crops[7].

As discussed in citegollin2019 and Moscona (2018), over the course of the Green Revolution, starting in the early 1960's, high-yield varieties were developed for different crops at different points in time, and tended to quickly diffuse across the world. Hence, as soon as a new HYV was discovered, it tended to be available for use in every country in the world more or less immediately, but countries varied in how useful this was depending on the climactic suitability of that country for that crop. One might be worried that the research and development of HYVs was itself targeted to crops that grew particularly well in particular countries, but this was largely not the case. For the most part, the Green Revolution was the result of highly centralized global philanthropic research efforts rather than country-specific profit-motivated R&D. Furthermore, note in equation (5) that $P_{i,t}$ is an *unweighted* average across crops, hence it is not affected by the endogenous selection of crops in different countries in either the high- or low- productivity states.

Table 2 shows the results of our IV estimation for both of our outcomes of interest. First stage results are available in Appendix A.

We find a strong positive association between both openness and productivity for both measures of the non-agricultural share and a strong negative association between their interaction and the non-agricultural share of GDP and the labor force.[8]

[7] Moscona (2018) argues that the FAO data are structured such that the low- versus high-input potential yields correspond to the HYV technology difference.

[8] For our IV results, we use robust standard errors, but do not cluster. For OLS, we cluster at the country level as there is reason to believe errors in the level of exports or productivity are serially correlated by country, however, the country-specific components of our instruments are static (i.e. the distance from other countries and the agro-climatic suitability of crops) which should be accounted for entirely by country-fixed effects, while the time-variation comes from changes in other-country GDP and HYV availability which is correlated across rather than within countries.

	(1) Non-Ag %GDP	(2) Non-Ag %LF
Trade/GDP	0.381	0.829*
	(0.289)	(0.428)
ProdCal	0.455**	1.551***
	(0.209)	(0.390)
Trade/GDP X ProdCal	-0.009**	-0.023***
	(0.004)	(0.005)
Year FE	Yes	Yes
Country FE	Yes	Yes
N Obs	4124	3979

Table 2: IV

2.4 Total Effect of Productivity

To examine the relevance of this finding, we briefly examine whether there are countries with a total negative effect on the non-agricultural share from an increase in agricultural productivity. Note that, in this reduced-form analysis, whether the total effect is positive or negative is simply a matter of whether the country is above or below a cutoff value of openness.

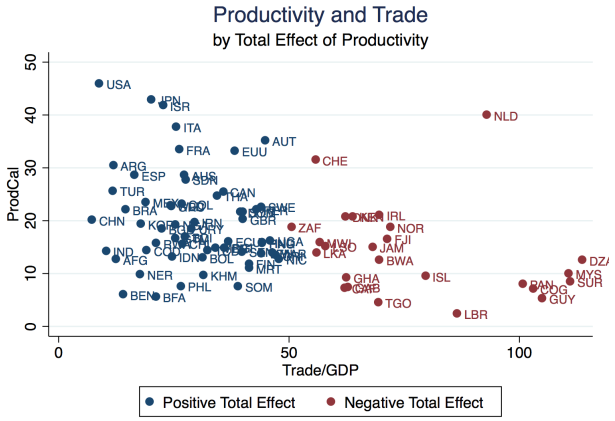


Figure 1: Countries by total effect of productivity in 1961

As Figure 1 shows, in 1961 a handful of countries actually have a negative total effect of increasing agricultural productivity on their non-agricultural share. Moreover, note that the share of countries for which the expected total effect of productivity is negative is increasing over time, due to trade openness increasing over time across the world, as seen in Appendix A.

3 Model

Our reduced-form analysis shows that the effect of agricultural productivity growth on structural change depends crucially on a country's openness to trade. However, one limitation of the reduced-form analysis is that it treats each country as a separate observation. That is, it supposes that the effects of agricultural productivity and openness on a country's structural change are simply a matter of the amount of agricultural productivity growth in that country itself and the country's overall level of openness.

In reality, it potentially matters not only how much agricultural productivity is growing in that country itself but also in other countries, and it potentially matters not only how open to trade that country is overall, but also which specific countries it is trading with. The only way to explore this intricate web of effects is through a structural model. To that end, we turn our attention now to a structural model of international trade and structural change. In this section we describe the model, then in section 4 we simulate the model to show how the effects of agricultural productivity growth on sectoral labor shares vary with openness. In section 5 we take the model to data in order to quantify the effects of the Green Revolution on each country's structural change, depending on where each country is in the global trading network.

Our model is a multi-country, two-sector (agriculture and non-agriculture) Ricardian model of trade. The structure of the model follows Eaton and Kortum (2002) ("EK"), but with non-homothetic CES preferences over two sectors, following Comin et al. (2017) ("CLM").[9]

[9]Note that CLM's model is a closed-economy model; we embed it in the EK model of trade. Our model is also similar to Uy et al. (2013) ("UYZ"), except that preferences in UYZ are such that the extent of non-homotheticity varies substantially by income level (in particular, as income goes to infinity, preferences are asymptotically homothetic). CLM find evidence in both international data and sub-national data that the extent of non-homotheticity is roughly constant across income levels, which their preference structure rationalizes. Given that our analysis looks at a much broader range of countries than UYZ, this is important for our purposes.

3.1 Production and trade

There is a continuum of goods in the agriculture (A) and non-agriculture (M) sectors. Each country has the technology to produce all goods in all sectors, but with different levels of productivity. The production function for good $z \in [0,1]$ in sector $k \in \{A, M\}$ in country i is

$$Y_{ik}(z) = A_{ik}(z)L_{ik}(z) \tag{6}$$

where $Y_{ik}(z)$ denotes output, $A_{ik}(z)$ denotes exogenous productivity, and $L_{ik}(z)$ denotes labor.

Following EK, $A_{ik}(z)$ is the realization of a random variable Z_{ik}, with each $A_{ik}(z)$ drawn independently from a Fréchet distribution with cumulative distribution function $F_{ik} \equiv Pr[Z_{ik} \leq A] = e^{-T_{ik}A^{-\theta}}$, where $T_{ik} > 0$ and $\theta > 1$. The larger T_{ik} is, the higher is country i's mean productivity in sector k. The higher θ is, the less dispersion there is in productivity across goods within a sector within a country.

Trade between countries is subject to iceberg trade costs. That is, if one unit of good z in sector k is shipped from country j to country i, then only $1/\tau_{ijk}$ units arrive in country i, where $\tau_{ijk} \geq 1$. We assume zero trade costs within a country, i.e., $\tau_{iik} = 1$ $\forall i, k$.

Markets are perfectly competitive, and prices are therefore determined by marginal costs of production. The price at which country j can supply good z in sector k to country i is $p_{ijk}(z) = \tau_{ijk}w_j/A_{jk}(z)$, where w_j is the wage in country j. Consumers buy each good from the cheapest source, hence the price of good z in sector k in country i is $p_{ik}(z) = \min\{p_{ijk}(z)\}$.

The composite good in sector k consumed by country i, denoted C_{ik}, is a CES aggregate of the individual goods $C_{ik}(z)$:

$$C_{ik} = \left(\int_0^1 C_{ik}(z)^{\frac{\eta-1}{\eta}} dz\right)^{\frac{\eta}{\eta-1}} \tag{7}$$

where $\eta > 0$ is the elasticity of substitution across goods within a sector.

Given the Fréchet distribution of productivities, EK show that the price of the composite good C_{ik} is $P_{ik} = \Gamma(\Phi_{ik})^{-\frac{1}{\theta}}$, where the constant Γ denotes the Gamma function evaluated at $(1 - \frac{\eta-1}{\theta})^{\frac{1}{1-\eta}}$, and $\Phi_{ik} = \sum_j T_{jk}(w_j \tau_{ijk})^{-\theta}$ captures country i's access to global production in sector k.[10]

The share of country i's expenditure on sector k sourced by country j equals the probability that country j is the lowest-cost supplier of a good z in sector k to country i. This equals

$$\pi_{ijk} = \frac{T_{jk}(w_j \tau_{ijk})^{-\theta}}{\Phi_{ik}} \tag{8}$$

3.2 Preferences

Following CLM, the representative household in country i has non-homothetic CES utility U_i implicitly defined by the equation

$$\sum_k (\Omega_k U_i^{\epsilon_k})^{\frac{1}{\sigma}} C_{ik}^{\frac{\sigma-1}{\sigma}} = 1 \tag{9}$$

where σ is the elasticity of substitution between the two sectors, Ω_k is a sector-k utility weight, and ϵ_k governs the income elasticity of demand for sector k. As utility U_i rises, the weight given to the consumption of sector k varies at a rate controlled by ϵ_k. More specifically,

$$\frac{\partial \log(C_{iA}/C_{iM})}{\partial \log U_i} = \epsilon_A - \epsilon_M \tag{10}$$

As CLM show, Ω_k and ϵ_k can be normalized to 1 for one sector without loss of generality; we normalize Ω_M and ϵ_M to 1.[11]

The representative household is endowed with a unit of labor, which it supplies inelastically, and it spends all of its labor income. It maximizes its utility U_i in (9) subject

[10] Under the assumption $\eta - 1 < \theta$, this price index is well-defined, and the elasticity of substitution η can be ignored because it only appears in the constant term Γ.

[11] See CLM for more discussion of non-homothetic CES preferences, in particular for a characterization of the Hicksian and Marshallian demand functions.

to its budget constraint

$$\sum_k P_{ik} C_{ik} = w_i \tag{11}$$

3.3 Equilibrium

Labor is perfectly mobile across sectors within a country, but immobile across countries. The labor market clearing condition in country i is

$$\sum_k L_{ik} = L_i \tag{12}$$

where L_i is the total labor endowment in country i, and L_{ik} is the labor employed in country i in sector k.

Trade between countries is balanced:

$$w_i L_i = \sum_j \sum_k \pi_{jik} \omega_{jk} w_j L_j \tag{13}$$

where $\omega_{jk} \equiv \frac{P_{jk} C_{jk}}{\sum_k P_{jk} C_{jk}}$ is the share of country j's expenditure on sector k.

Definition of equilibrium. Given a set of parameters $\{L_i, T_{ik}, \theta, \tau_{ijk}, \eta, \Omega_k, \epsilon_k, \sigma\}_{i,j,k}$, an equilibrium of the economy is a set of prices $\{P_{ik}, w_i\}_{i,k}$, allocations $\{Y_{ik}, L_{ik}, C_{ik}\}$, and trade shares $\{\pi_{ijk}\}_{i,j,k}$ such that: (i) given prices, the allocations solve the firms' profit-maximization problem associated with technologies (6) and households' utility-maximization problem given by (9) and (11); and (ii) the allocations satisfy the market clearing conditions (12) and (13), with trade shares given by (8).

4 Numerical illustration

Suppose a country has an increase in its agricultural productivity. What effect does this have on structural change − in particular, does this increase or decrease the share of the country's labor force working in the agricultural sector? Our model captures

two opposing effects. On the hand, higher agricultural productivity raises the country's income, which shifts demand toward non-manufacturing, which thereby causes a *decrease* in the equilibrium share of the country's labor force in agriculture. This is the *labor push effect*. On the other hand, higher agricultural productivity means the country now has a larger comparative advantage in agriculture, which causes an *increase* in the country's equilibrium share of the labor force in agriculture. The more open the country's economy is, the stronger the latter effect is relative to the former.

To help illustrate this, in this section we perform a numerical simulation of the model. For the sake of simplicity, in this section we use a stripped-down version of the model in section 3, with just two countries and with a simpler structure of non-homothetic preferences. In particular, the representative household in country i's utility U_i is Stone-Geary:

$$U_i = (C_{iA} - \overline{C_A})^{\beta_A}(C_{iM})^{1-\beta_A} \tag{14}$$

where $\overline{C_A}$ denotes an agricultural subsistence requirement. Utility is only positive for $C_{iA} > \overline{C_A}$, hence the household first spends all of its income on agricultural goods, and then spends β_A fraction of its remaining income on agriculture, $1 - \beta_A$ on non-agriculture. Hence the income elasticity of demand is greater than one for non-agriculture and less than one for agriculture.

The parameter values we use in our simulation simulation are given in Table 3. The two countries are identical in every way, except for the parameter T_{iA} governing agricultural productivity, which we vary from 5 to 20 for country 1 (while holding it fixed at 10 for country 2, and also holding T_{iM} fixed at 10 for both countries).

For simplicity, trade costs here are symmetric. That is, $\tau_{12A} = \tau_{21A} = \tau_{12M} = \tau_{21M} \equiv \tau$.[12] At the same time that we vary T_{1A} from 5 to 20, we also vary τ from 1 to 2. Note that $\tau = 1$ represents free trade between the two countries, while $\tau = 2$ represents trade costs of 100%.

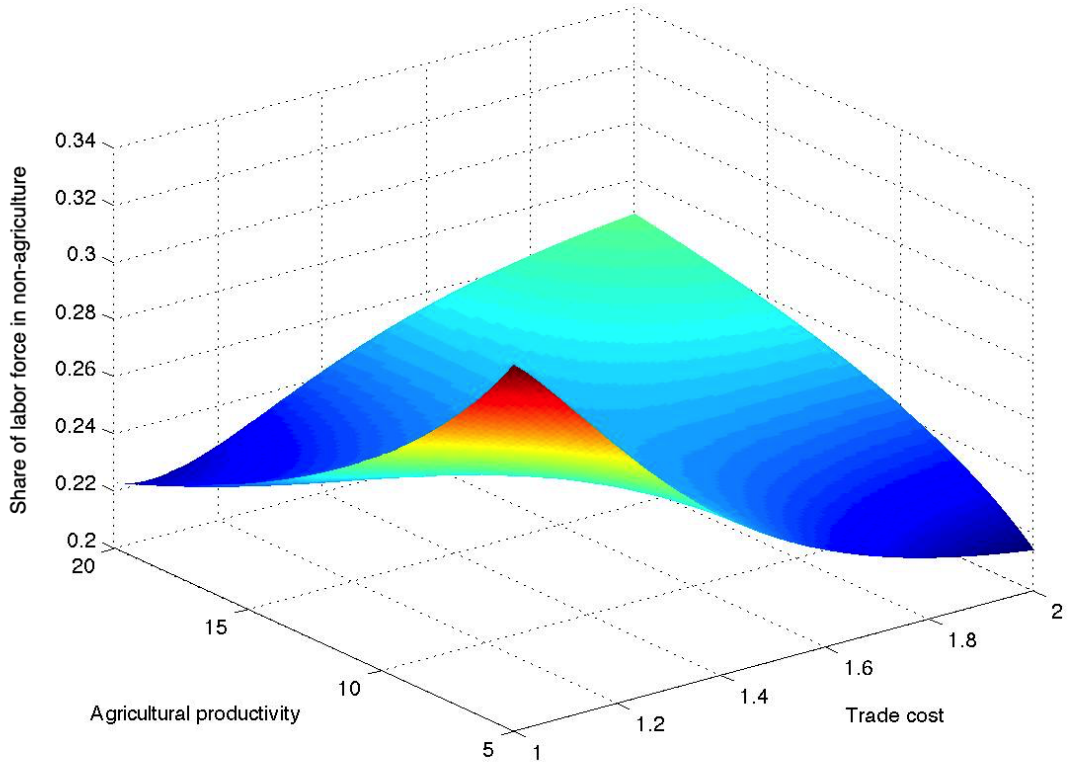
[12]Note, though, that there is free trade within countries, i.e., $\tau_{iiA} = \tau_{iiM} = 1$ for $i = 1, 2$.

Table 3: Parameter values for numerical illustration

Parameter	Meaning	Value
L_i	Size of labor force in each country	1
θ	Within-sector productivity dispersion	5
η	Within-sector elasticity of substitution	3
$\overline{C}_A$	Agricultural subsistence requirement	1
β_A	Agricultural weight in utility function	0.5
T_{iM}	Non-ag. productivity in each country	10
T_{2A}	Agricultural productivity in country 2	10
T_{1A}	Agricultural productivity in country 1	vary from 5 to 20
τ	Iceberg trade cost	vary from 1 to 2

For each combination of T_{iA} and τ, we simulate the model and compute the equilibrium share of country 1's labor force in non-agriculture. The results are shown in Figure 2. The vertical axis is country 1's equilibrium share of its labor force in non-agriculture. The two horizontal axes are T_{1A} and τ. Notice that at high levels of openness (e.g., $\tau = 1$), higher agricultural productivity in country 1 *decreases* country 1's non-agricultural labor share; this is the *comparative advantage effect*. Meanwhile, at low levels of openness (e.g., $\tau = 2$), higher agricultural productivity in country 1 *increases* country 1's non-agricultural labor share; this is the *labor push effect*. One can also see in the figure that, between the two extremes of $\tau = 1$ and $\tau = 2$, the effect varies continuously.

Figure 2: Numerical illustration of the labor push vs. comparative advantage effect



Matsuyama (1992) analyzed the two extreme cases of a closed economy (which in this model corresponds to $\tau \to \infty$) and a small open economy (which in this model corresponds to $\tau = 1$ and $L_1 \to 0$). What we have shown in this section, using an Eaton-Kortum model of trade with non-homothetic preferences, is that Matsuyama's point holds under cases of intermediate levels of openness as well, with the relative strength of the labor push effect and the comparative advantage effect varying in a continuous way with openness.

This numerical illustration shows this in the simplest, starkest possible way, but given that the parameter values in Table 3 are arbitrary, this raises the question of how strong the labor push and comparative advantage effects actually are. In the next section we take the model of section 3 to data in order to answer this question.

5 Quantitative analysis

In this section we examine plausibly exogenous changes in countries' agricultural productivity during the Green Revolution. The previous section showed, through a simple two-country numerical simulation, that whether an increase in agricultural productivity accelerates or retards a country's structural change (more precisely, whether it increases or decreases the share of the country's labor force in agriculture) depends on the level of openness of the country to the rest of the world. During the Green Revolution, when many different countries were simultaneously experiencing increases in agricultural productivity – with the size of this increase varying from country to country – and these countries were trading with each other subject to a complex network of trade costs varying at the country-pair level, the model in section 3 makes a rich set of predictions regarding how the Green Revolution changed the global trading equilibrium, in particular how much each country's sectoral labor shares changed.

In this section we evaluate these predictions. In particular, we calibrate the model from section 3 using 1961 data, then feed into the model the plausibly exogenous changes in agricultural productivity that occurred from 1961 to 1995 during the Green Revolution, and then compute the new equilibrium of the model. We show that the model is able to match much of the observed changes from 1961 to 1995 in countries' sectoral labor shares.

5.1 Calibration

The data we use for our structural analysis span from 1961 to 2016 and cover 67 countries.[13]

L_i is calibrated for each country i to match the country's 1961 population according to the World Bank's World Development Indicators (WDI).

The parameter θ that governs the amount of within-country, within-sector productivity dispersion is calibrated to a value of 4, based on estimates from Simonovska and

[13] More than 200 countries are in the raw data. We dropped countries for which not all the variables used in this section were available for 1961.

Waugh (2014). The within-sector elasticity of substitution η is set to 3.[14] The elasticity of substitution σ between agriculture and non-agriculture is set to 0.35, based on estimates from CLM.

For trade costs τ_{ij} between each pair of countries, we use estimates of these costs from the ESCAP-World Bank Trade Cost Database from Arvis et al (2013). Their estimates that we use are at the country-pair level, rather than sector-specific; we assume $\tau_{ijA} = \tau_{ijM} \equiv \tau_{ij}$. These estimates are available only for the years 1995 to 2016, and for many country-pairs they are missing during certain years, so we average across years for each country-pair. This means we are assuming that, for each pair of countries i and j, τ_{ij} was the same in 1961 as it was in 1995-2016; this is admittedly a crude assumption, but there is no sense in which it necessarily biases our results in a particular direction.

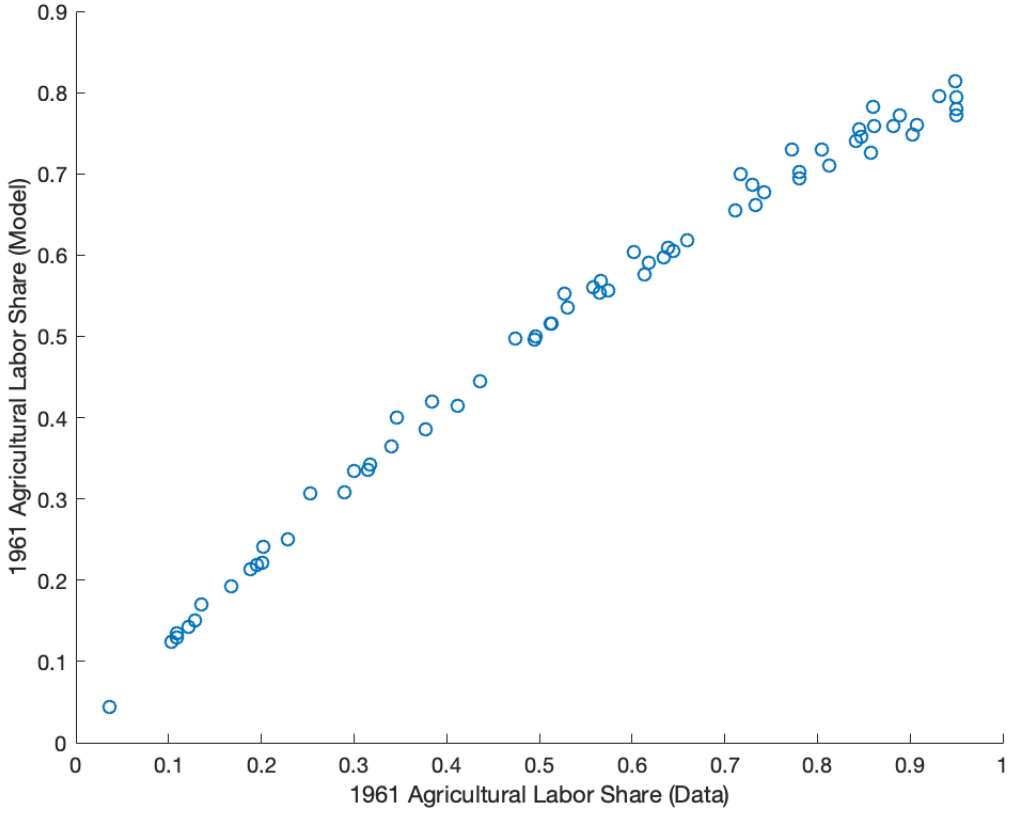
T_{iA} is calibrated for each country i based on its 1961 productivity in agriculture. Following Moscona (2018), agricultural productivity is measured in terms of calories per acre, using data from the FAO on each country's production of each crop (rice, wheat, maize, sorghum, potato, groundnut, dry beans, and millet), combined with estimates from the USDA of how many calories per kilogram there are in each of these crops.

$\{T_{iM}\}_i$ (each country's non-agricultural productivity), Ω_A (the weight on agriculture in the utility function, relative to non-agriculture), and ϵ_A (the income elasticity of demand for agricultural goods, relative to non-agricultural goods) are calibrated to match countries' 1961 agricultural labor shares, the data for which are borrowed from Wingender (2014). We estimate Ω_A to be 0.8 and ϵ_A to be 0.9.

Figure 3 depicts each country's 1961 share of employment in agriculture. The horizontal axis is the share observed in the data, while the vertical axis is the share predicted by the model. Each circle in the figure is a country. Note that these are targeted moments. The model does a very good job of matching these for most countries, but somewhat under-predicts agricultural labor shares for the countries whose agricultural share of employment is over 0.7.

[14] Recall from footnote 10 that η does not actually affect any of the equilibrium objects of interest in the model.

Figure 3: 1961 agricultural labor shares, model vs. data



5.2 Model validation: the "Green Revolution shock"

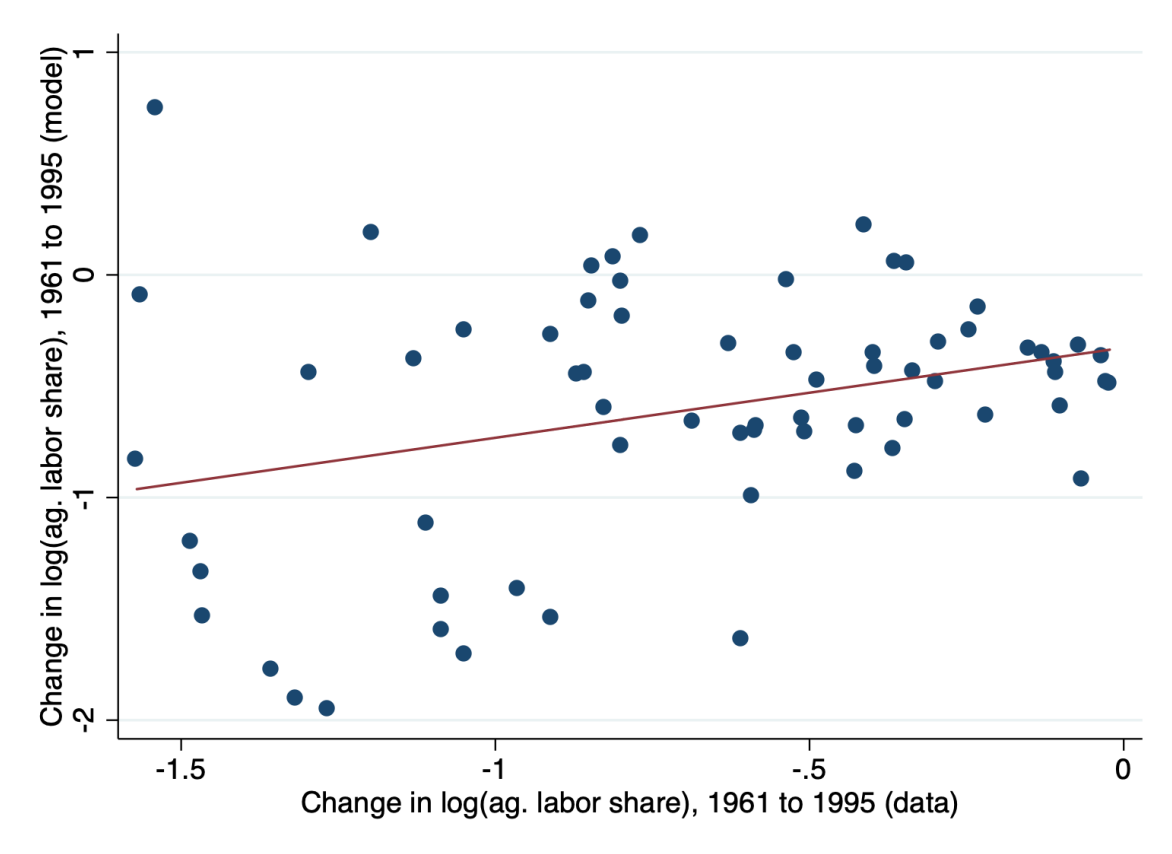
With the model from section 3 calibrated as described in section 5.1, in this section we feed into the model the "Green Revolution shock." That is, we change T_{iA} for each country, based on the plausibly exogenous change from 1961 to 1995 in the Moscona (2018) predicted productivity. Holding all of the other parameters of the model fixed, we then re-compute the equilibrium of the model under these new values of T_{iA}. In particular, we calculate the model's prediction for how much each country's agricultural labor share changes from 1961 to 1995, and compare this against the data.

As emphasized throughout our discussion of the model, an important source of expected variation in the effects of agricultural productivity growth on structural change is countries' differing levels of openness. Consider, for example, the case of two African

countries in our sample, Benin and Burundi. In 1961, both of these countries had very large shares of employment in agriculture: 84.5% in Benin, 95% in Burundi. (Our structural model predicts these 1961 shares to be 75.5% for Benin and 78% for Burundi.) Between 1961 and 1995, these two countries experienced almost identical shocks to their predicted agricultural productivity: an increase of 256% for Benin, 266% for Burundi. However, these two countries had very different levels of openness. In 1961, Benin's exports-to-GDP ratio was only 0.046, while Burundi's was 0.13. We should therefore expect Benin's agricultural labor share to decrease significantly more than Burundi's – our structural model predicts Benin's agricultural labor share to decrease to 38.3% by 1995, but only decrease to 48.3% for Burundi. And indeed that is what happened in the data, even more extremely: from 1961 to 1995 Benin's agricultural labor share decreased to 47.1%, while Burundi's only decreased to 92.5%.

Going beyond the specific example of Benin and Burundi, the results for the full set of 67 countries are shown in Figure 4. The horizontal axis is the log change from 1961 to 1995 in countries' agricultural share of employment observed in the data. The vertical axis is the change predicted by the model.

Figure 4: 1961-1995 changes in agricultural labor shares, model vs. data



Note in Figure 4 that for 8 of the 67 countries (8 countries with mostly high levels of openness), the structural model predicts that their structural change was actually impeded by the Green Revolution. This is 12% of the sample, which roughly lines up with the 18% of the sample for which a negative effect was predicted by our reduced-form analysis.

Regressing the model's predictions on the data gives an estimated coefficient of 0.40, with a standard error of 0.15, a p-value of 0.009, and an R-squared of 0.10. This significantly positive fit between the model's predictions and the data is remarkable, given that in this exercise we are using absolutely no information on any changes in any variables between 1961 and 1995 other than the Green-Revolution-induced changes in agricultural productivity. This implies that the Green Revolution was a significant driver of changes in sectoral labor shares during this time period, and that our structural model

does a good job of explaining the heterogeneity across countries in the effects of the Green Revolution on structural change.

6 Final comments

Among both academic economists and policymakers, it is often taken as a given that growth in agricultural productivity has purely beneficial effects on structural change and economic growth. Ever since Matsuyama (1992) raised the possibility that the effects could be negative in open economies, this point has been well-appreciated by economists at a theoretical level, but has typically been brushed aside at a practical level, due to his small open economy assumption being so extreme. In this paper we caution against this brushing aside. We show through a two-sector Eaton-Kortum model of trade that the effects can be negative even at intermediate (and realistic) levels of openness. We back this up with reduced-form evidence, exploiting the Green Revolution as a plausibly exogenous source of changes in agricultural productivity.

There is room for future research in several directions. Our analysis is at the country level. Given the recent explosion of research using structural quantitative models in economic geography, our analysis could be carried out within countries. One might expect the comparative advantage effect to dominate even more at the level of localities than it does at the level of countries.

One limitation of our analysis is a lack of information on the labor intensity of the new agricultural technologies whose diffusion during the Green Revolution we exploit (accordingly, in our model we assume for simplicity that labor is the only factor of production). In contrast, Bustos et al. (2016) carefully examine the effects on structural change of the introduction of new agricultural technologies with differing levels of labor intensity, but do not examine how the effects vary with openness. An analysis that simultaneously examines both of these dimensions has, to our knowledge, not yet been done and would be valuable.

One of the motivations for examining structural change is its implications for eco-

nomic growth. For example, in the theoretical model of Matsuyama (1992), the non-agricultural sector experiences endogenous growth from learning-by-doing while the agricultural sector does not, hence movement of labor into the non-agricultural sector causes an increase in aggregate economic growth, while the pulling of labor into agriculture retards it. Given the additional identification challenges inherent in exploring such a thing, we hold off on making any claims in that direction in this paper, but it is a direction that is begging to be further explored both theoretically and empirically.

References

Alvarez-Cuadrado, F. and Poschke, M. (2011). Structural change out of agriculture: Labor push versus labor pull, *American Economic Journal: Macroeconomics* pp. 127–158.

Bustos, P., Caprettini, B. and Ponticelli, J. (2016). Agricultural productivity and structural transformation: Evidence from brazil, *American Economic Review* **106**(6): 1320–1365.

Comin, D., Lashkari, D. and Mestieri, M. (2017). Structural change with long-run income and price effects.

Eaton, J. and Kortum, S. (2002). Technology, geography, and trade, *Econometrica* **70**(5): 1741–1779.

Feyrer, J. (2019). Trade and income—exploiting time series in geography, *American Economic Journal: Applied Economics* **11**(4): 1–35.

Gollin, D., Parente, S. L. and Rogerson, R. (2007). The food problem and the evolution of international income levels, *Journal of Monetary Economics* **54**(4): 1230–1255.

Gollin, D., Parente, S. and Rogerson, R. (2002). The role of agriculture in development, *American economic review* **92**(2): 160–164.

Helliwell, J. F. (2000). *How much do national borders matter?*, Brookings Institution Press.

Herrendorf, B. M., Rogerson, R. and Valentinyi, Á. (2014). Growth and structural transformation, *Handbook of Economic Growth* **2**: 855–941.

Matsuyama, K. (1992). Agricultural productivity, comparative advantage, and economic growth, *Journal of Economic Theory* **58**(2): 317–334.

Moscona, J. (2018). The impact of india's green revolution: An empirical investigation of modern agricultural development, *Unpublished Manuscript* .

Ngai, L. R. and Pissarides, C. A. (2007). Structural change in a multisector model of growth, *American economic review* **97**(1): 429–443.

Nurkse, R. (1953). *Problems of capital formation in underdeveloped countries*, Oxford University Press.

Rostow, W. W. (1960). *The stages of growth: A non-communist manifesto*, Cambridge University Press Cambridge.

Schultz, T. W. (1953). *The economic organization of agriculture*, McGraw-Hil.

Uy, T., Yi, K.-M. and Zhang, J. (2013). Structural change in an open economy, *Journal of Monetary Economics* **60**(6): 667–682.

Wingender, A. M. (2014). Structural transformation in the 20th century: A new database on agricultural employment around the world.

A Appendix

A.1 Alternate measures of trade

Table 4 shows the OLS results replacing export share with total trade share (Columns 1 and 2) and import share (Columns 3 and 4).

	(1) Non-Ag %GDP	(2) Non-Ag %LF	(3) Non-Ag %GDP	(4) Non-Ag %LF
ProdCal	0.151**	0.058	0.003	0.014
	(0.064)	(0.070)	(0.063)	(0.077)
Exp/GDP	0.192***	0.057		
	(0.063)	(0.051)		
Exp/GDP X ProdCal	-0.004***	-0.002		
	(0.001)	(0.002)		
Imp/GDP			0.020	0.021
			(0.047)	(0.031)
Imp/GDP X ProdCal			0.000	-0.001
			(0.001)	(0.001)
Year FE	Yes	Yes	Yes	Yes
Country FE	Yes	Yes	Yes	Yes
Cluster	Country	Country	Country	Country
N Obs	6339	5817	6339	5817

Table 4: OLS with alternate trade definitions

Table 5 shows the results using the Moscona IV and trade distance-weighted partner GDP strategies.

	(1) Non-Ag %GDP	(2) Non-Ag %LF	(3) Non-Ag %GDP	(4) Non-Ag %LF
ProdCal	0.228**	1.233***	0.813	1.455***
	(0.100)	(0.283)	(1.450)	(0.524)
Exp/GDP	0.349*	1.089**		
	(0.183)	(0.475)		
Exp/GDP X ProdCal	-0.013***	-0.039***		
	(0.003)	(0.008)		
Imp/GDP			4.427	3.255
			(12.033)	(3.804)
Imp/GDP X ProdCal			-0.040	-0.034
			(0.077)	(0.022)
Year FE	Yes	Yes	Yes	Yes
Cluster	Yes	Yes	Yes	Yes
N Obs	4124	3979	4124	3979

Table 5: IV with alternate trade definitions

Consistent with the underlying theory, the results are slightly stronger for exports (where comparative advantage and openness interact directly) than for imports, though are findings are similar for both.

A.2 Restricting to non-OECD countries

	(1) Non-Ag %GDP	(2) Non-Ag %LF
Trade/GDP	7.274	1.961
	(21.384)	(2.044)
ProdCal	18.404	6.555
	(54.623)	(6.502)
Trade/GDP X ProdCal	-0.212	-0.064
	(0.618)	(0.061)
Year FE	Yes	Yes
Country FE	Yes	Yes
N Obs	3280	2862

Table 6: IV, non-OECD countries

Table 6 displays the IV results excluding OECD countries. Though we lose significance, the coefficients are similar to the main results with a positive association between trade and productivity with non-agricultural share and a negative association between the interaction term and non-agricultural share for both measures.

A.3 Alternate distance measure for Trade IV

We consider alternate measures of distance in constructing the Trade IV, squaring distance in the denominator and raising it to the 4th.

A.4 Robustness to functional form

We consider the robustness of the interaction term by including de-meaning the interaction term and including squared terms for the main regressors.

Each of the combinations of the main analysis for both non-agricultural GDP share and non-agricultural labor force share hold for de-meaning and adding a squared term, adding confidence that the main result is robust to misspecification.

	(1) Non-Ag %GDP	(2) Non-Ag %LF	(3) Non-Ag %GDP	(4) Non-Ag %LF
Trade/GDP	0.325	0.821***	0.465**	0.803
	(0.202)	(0.241)	(0.228)	(0.571)
ProdCal	0.336***	1.606***	0.817**	1.054
	(0.098)	(0.270)	(0.406)	(1.378)
Trade/GDP X ProdCal	-0.007***	-0.025***	-0.020	-0.003
	(0.002)	(0.004)	(0.016)	(0.034)
Distance Power	2	2	4	4
Year FE	Yes	Yes	Yes	Yes
Country FE	Yes	Yes	Yes	Yes
N Obs	4124	3979	4124	3979

Table 7: IV with distance definitions

	(1) Non-Ag %GDP	(2) Non-Ag %GDP	(3) Non-Ag %GDP	(4) Non-Ag %LF	(5) Non-Ag %LF	(6) Non-Ag %LF
Trade/GDP X ProdCal	-0.013*			-0.025**		
	(0.008)			(0.012)		
Trade/GDP X ProdCal, de-meaned		-0.009**	-0.013*		-0.023***	-0.025**
		(0.004)	(0.008)		(0.005)	(0.012)
Trade/GDP	1.076	0.110	0.708	0.204	0.179	-0.507
	(0.875)	(0.187)	(0.676)	(1.466)	(0.316)	(1.171)
ProdCal	-0.836	-0.211*	-1.741	0.465	-0.047	-1.284
	(1.250)	(0.122)	(1.748)	(2.314)	(0.106)	(3.086)
Squared Terms	Yes	No	Yes	Yes	No	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Country FE	Yes	Yes	Yes	Yes	Yes	Yes
N Obs	4124	4124	4124	3979	3979	3979

Table 8: Robustness to functional form

A.5 First stage results

Table 9 gives the first stage results for the IV analysis. Each instrument is well predicted in total, though when including the interacted effects, the signs are not always as expected.

As 11 shows, the instruments for trade share and caloric yield well predict the value of each individually.

We also check the robustness of the IV by including five year leads as predictors for the variable of interest.

We find the same-year variable to be a stronger predictor for caloric production, however, we find the lead term to be a stronger predictor for trade. This is of some concern in that it is indicative of the persistence of our IV for openness, though there

	(1)	(2)	(3)
	Trade/GDP	ProdCal	Trade/GDP X ProdCal
dGDP	0.001	-1.586***	1e-4**
	(0.001)	(0.484)	(6e-5)
P	-0.003***	-0.224	3e-4***
	(0.001)	(0.205)	(3e-5)
dGDP X P	-0.030	-63.21***	0.024***
	(0.080)	(20.19)	(0.003)
Year FE	Yes	Yes	Yes
Country FE	Yes	Yes	Yes

Table 9: First stage

	(1)	(2)	(3)
	Trade/GDP	ProdCal	Trade/GDP X ProdCal
dGDP	0.003***		
	(0.000)		
P		0.960***	
		(0.124)	
dGDP X P			-0.008
			(0.007)
Year FE	Yes	Yes	Yes
Country FE	Yes	Yes	Yes

Table 10: First stage, excluding other regressors

	(1)	(2)
	Trade/GDP	ProdCal
dGDP, 5 year lead	0.007***	
	(0.002)	
dGDP	-0.005***	
	(0.002)	
P, 5 year lead		-0.920**
		(0.361)
P		1.213***
		(0.172)
Year FE	Yes	Yes
Country FE	Yes	Yes

Table 11: First stage with leads

is also decent reason to believe that most of the changes to the interaction should come from changes in productivity.

A.6 Countries with negative productivity effects

Expanding on Table **??**, we look at the share of countries that experience a negative total effect of productivity non-agricultural share of GDP and the labor force.

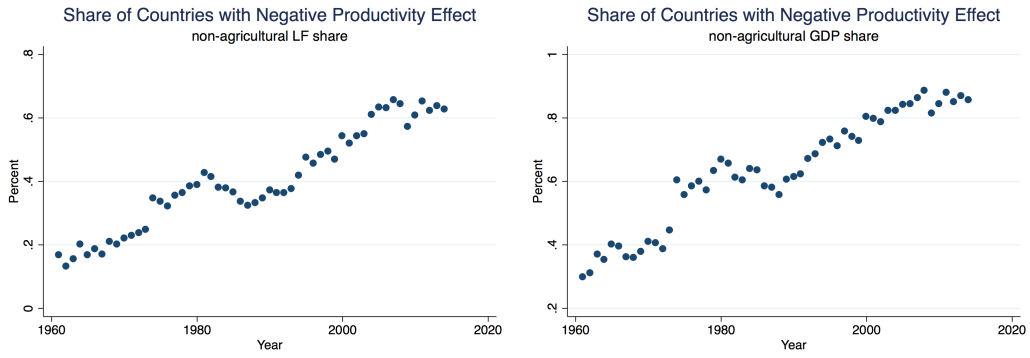


Figure 5: Lowess-smoothed binned results

As Figure 5 shows, as exports increase over time, the implied share of countries that have a negative total effect of productivity on their non-agricultural GDP and LF shares is increasing over time, indicating increasing relevance of this interaction.

A.7 Non-parametric estimates

Recognizing that the interaction between openness and productivity may be non-linear, we estimate a lowess-smoothed regression and plot the residuals of the interaction term and the non-agricultural share of GDP from the main regression with and without country fixed effects.

As Figure 6 shows, the relationship between the interaction between the exports-productivity interaction is quite linear. This adds confidence to the linear specification of our empirical estimation.

A.8 Year by Year Coefficients

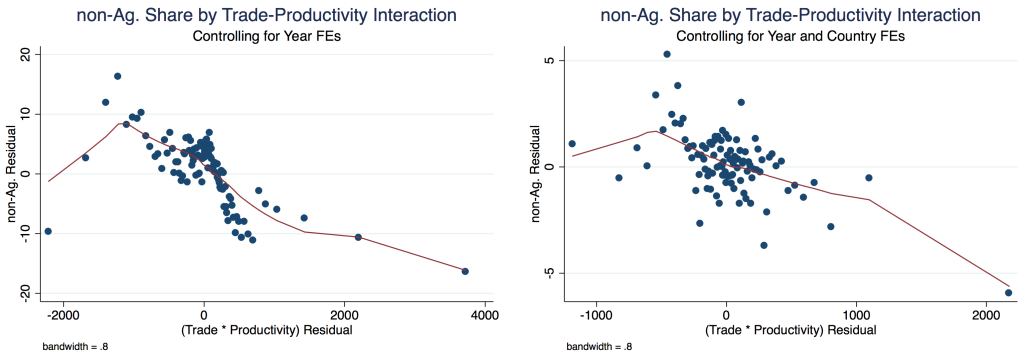


Figure 6: Lowess-smoothed binned results

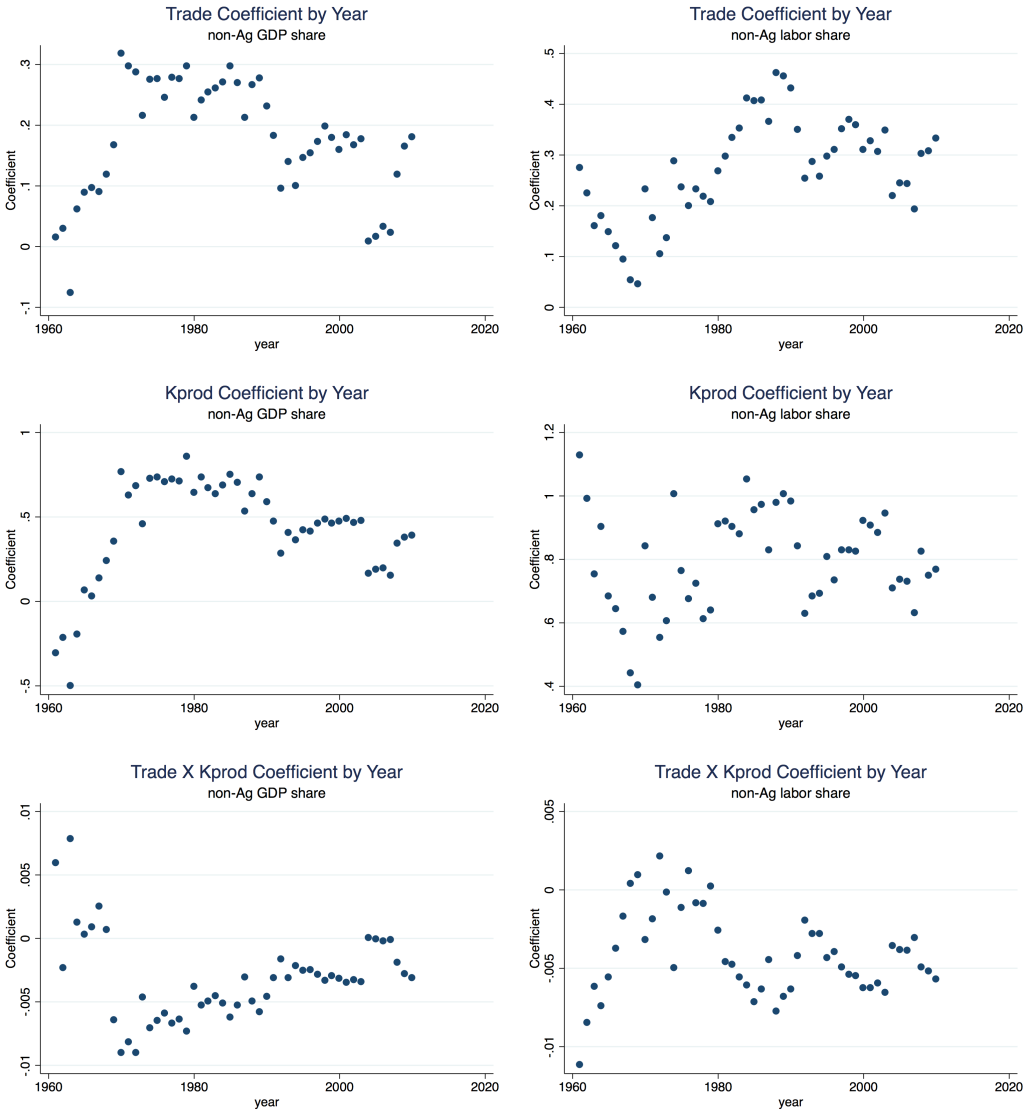


Figure 7: Coefficients bye year