Specialization, Economic Development and Aggregate Productivity Differences

David Lagakos

Arizona State University

Michael E. Waugh

New York University

November 2009

ABSTRACT

Cross-country labor productivity differences are large in agriculture and much smaller in non-agriculture. We argue that these relative productivity differences arise when subsistence consumption needs prevent workers in poor countries from specializing in the sector in which they are most productive. We formalize our theory by embedding the Roy (1951) model of selection into a two-sector general-equilibrium growth model in which the agents' preferences feature a subsistence food requirement. A parameterized version of the model predicts that output per worker gaps will be substantially larger across countries in agriculture than non-agriculture even though countries differ only by a sector-neutral efficiency term.

Email: lagakos@asu.edu, mwaugh@stern.nyu.edu. For helpful comments we thank Francesco Caselli, Berthold Herrendorf, Narayana Kocherlakota, Erzo Luttmer, Ellen McGrattan, Guillermo Ordoñez, Andrés Rodríguez-Clare, Richard Rogerson, Jim Schmitz and seminar participants at Arizona State, Clemson, the Federal Reserve Banks of Minneapolis and San Francisco, Iowa, Toronto, York, the 2009 Midwest Macro Meetings, the NBER Summer Institute Economic Growth group, the SED (Istanbul), and the SITE Growth and Development session. For sharing data with us we thank Tasso Adamopoulos and Keith Wiebe. All potential errors are our own.

1 Introduction

Cross-country labor productivity differences are large in agriculture and much smaller in non-agriculture relative to aggregate differences (Caselli, 2005; Restuccia, Yang and Zhu, 2008). Development accounting exercises have shown that these sector productivity differences are key in accounting for aggregate productivity differences. If agricultural labor productivity were hypothetically raised to the U.S. level in every country, or if the share of labor in agriculture were hypothetically lowered to the U.S. level, then international variation in aggregate productivity would be virtually eliminated (Caselli, 2005). These results suggest that understanding productivity differences in agriculture and non-agriculture are at the heart of understanding world income inequality.

In this paper we provide a theory of why labor productivity differences are larger in agriculture and smaller in non-agriculture than in the aggregate. We argue that these sector productivity differences arise when sector-neutral efficiency differences combine with subsistence food consumption needs to generate variation in the extent to which workers specialize in the sector where they are most productive.

The basic idea is that countries with low efficiency must deploy a large fraction of their workforce into the agriculture sector to satisfy subsistence food needs. As a result many of those working in agriculture are those whose comparative advantage is *not* in agricultural work, but rather in non-agricultural tasks such as writing newspaper articles, doing economic research, or teaching yoga classes. In countries with high efficiency, in contrast, a smaller fraction of workers are in agriculture, and those remaining in agriculture are those who are relatively most productive at farm work. As a result, physical productivity differences are larger in agriculture than in the aggregate. In non-agriculture the mechanism leads to exactly the opposite result, namely that productivity differences are smaller than in the aggregate.

We formalize our theory by embedding the Roy (1951) model of selection into a simple two-sector general-equilibrium growth model. Our theory has two main ingredients. First, workers are heterogenous in their ability to produce output in the two sectors and choose where to work. Second, preferences have a subsistence food requirement. Countries differ only in a sector-neutral efficiency term; preferences and the distribution of ability are identical across countries. Qualitatively, the model can generate productivity differences in agriculture that are larger than aggregate differences, and non-agriculture productivity differences that are smaller than in the aggregate. The novel feature of this result is it follows from optimal behavior only, as opposed to exogenous country-specific sectoral productivity differences, or barriers to agricultural production, as emphasized by other studies (e.g. Restuccia, Yang, Zhu, 2008).

Our main question of interest is whether the theory can quantitatively generate productivity differences that are substantially larger in agriculture than non-agriculture, as in the data. To answer this question we calibrate the model using parametric assumptions on the distribution of worker ability and observations on the distribution of wages in agriculture and non-agriculture in recent U.S. data. We show that the dispersion in ability for each sector is pinned down by the variances of wages for workers within the two sectors, while the correlation of the ability draws is disciplined by the ratio of average sector wages.

Our main exercise is to vary sector-neutral efficiency in the model so as to generate aggregate productivity differences equal to those found in the 90th and 10th percentile of the country income distribution, namely a factor of 22. We then compute the model's predictions for agriculture and non-agriculture productivity across countries and compare them to the data. If the model had no predictive power, it would predict that these sector productivity differences equalled those in the aggregate. We find that, to the contrary, the model generates a factor of 40 difference in agriculture productivity and a factor 10 difference in non-agriculture. In the data the corresponding sector productivity gaps are a factor of 45 in agriculture and 4 in non-agriculture. Thus, our model explains roughly three quarters of the difference between agriculture productivity gaps and aggregate gaps, and only slightly less of the difference between non-agriculture productivity gaps and those in the aggregate. We conclude that efficiency differences that affect each sector in the same way may nonetheless lead to agriculture differences that are much larger than those outside of agriculture.

We also show that the model performs quantitatively well in matching relevant development facts for the cross-section of counties. Specifically, the model is largely consistent the relationship between income per capita and the share of labor and GDP in agriculture, and performs moderately well in matching the relationship between income per capita and relative agricultural prices. We show that relative agricultural prices are higher in poor countries, with countries around the 10th percentile of the income distribution having prices around 2.5 times higher than countries in the 90th percentile. The model predicts this ratio should be around 4. We also compare our model's predictions to the same set of statistics for the U.S. time series, and show that the model's predictions are in line with the data. Finally, we show that the model's wage distribution closely resembles the current distribution of wages in United States.

To test the robustness of the model's predictive power, we ask what happens under a more conservative calibration of the model's most important two parameters, namely the variances of the two ability distributions. While this ability variance is disciplined in the benchmark model by cross-sectional wage variance in the United States, some economists have argued that some fraction of wage dispersion is unrelated to ability differences. Postel-Vinay and Robin (2002), for example, argue that around half of wage variation is due instead to labor market imperfections. To address this concern we recalibrate the model to have half the ability variance as the baseline calibration. In this more conservative calibration we still find that our model explains up to 50 percent of the productivity differences in agriculture and non-agriculture, relative to aggregate differences, between the 90th and 10th percentile of countries.

We conclude by providing direct evidence that our mechanism was at work in the development experiences of the United States and Britain. Two dimensions in which sector abilities are observable are sex and age: historians and development economists have argued that women and children have a comparative disadvantage in agricultural work relative to adult men. Our theory thus predicts that, during a structural transformation, women and children leave farm work and enter the industrial sector at a faster rate than men. We cite evidence that this is in fact what happened in Britain and the United States in the 18th and 19th centuries.

Our theory has new implications for the way economists think about aggregate productivity in the developing world. Concretely, our model suggests that low aggregate productivity is not *caused* by large fractions of workers working in the relatively unproductive agriculture sector. Instead, low measured productivity in agriculture and large agricultural labor shares are *consequences* of low sector-neutral productivity. The distinction is important because it helps determine the extent to which future research efforts on aggregate productivity differences should focus on the determinants of productivity in agriculture per se, as opposed to more general potential determinants. The policy implications between the two views are different as well. While accounting exercises suggest that fixing agriculture is crucial to raising overall productivity, our theory predicts that improvements in technology, institutions, or social infrastructure (Hall and Jones, 1999; Acemoglu, Robinson and Johnson, 2002) are the key to improving living standards.

2 Motivating Evidence and Related Literature

In this section we highlight the important role of agricultural in understanding aggregate productivity. Specifically, we reproduce the findings of Caselli (2005) to illustrate how differences in labor productivity and shares of workers in agriculture account for much of the variation in aggregate output per worker (as well as TFP) across countries. We then briefly discuss several other related studies which attempt to endogenize agriculture and non-agriculture sector productivity differences, as our paper does.

Panel A of Table 1 shows that labor productivity differences in agriculture are larger than aggregate differences, and that non-agriculture productivity differences are much smaller. The ratio of agricultural output per worker in the 90th and 10th percentiles of the income distribution is 45, compared to just 4 in non-agriculture. As a frame of reference, the ratio for the aggregates is 22. Panel B summarizes the well-known fact that poor countries have a much larger fraction of their work force in agriculture. A country whose per-capita income is in the 90th percentile

Table 1: **Agriculture and Labor-Productivity Accounting**

Panel A: Labor Productivity Differences

Sector	Ratio of 90th-10th Percentile
Agriculture	45
Aggregate	22
Non-Agriculture	4

Panel B: Percent of Labor in Agriculture

Country Income Percentile	Percent
90th	2.8
10th	78.3

Source: Caselli (2005)

has just 2.8% of its workers in agriculture, while the 10th percentile country has 78.3% of its workers in agriculture.

These two facts together highlight the potential importance of the agriculture/non-agriculture split. In an accounting sense, large aggregate differences in productivity are ”explained” by poor countries having virtually all their workers in the sector where productivity differences are largest relative to the richest countries. Caselli formalizes this argument by computing the hypothetical variance of cross-country aggregate output per worker assuming that agricultural productivity in all countries were equal to the U.S. level. His answer is just a factor of 1.6, down from the actual factor of 22! In other words, international labor productivity differences would be virtually eliminated. A similar experiment computes the hypothetical variance of aggregate output per worker assuming all countries had the U.S. share of workers in agriculture. This experiment yields a factor of 4.2 differences between the 90th and 10th percentile, which again is vastly lower than the 22 seen in the data.

One potential explanation for labor productivity differences in agriculture is physical capital per worker differences across countries. Caselli argues that labor productivity differences almost entirely represent total-factor productivity (TFP) differences. As he puts it, ”the factor-only model explains virtually nothing of the observed per-capita income variance in agriculture: it’s entirely a story of TFP differences, even more so than for aggregate GDP.” (Caselli, 2005, page 49.) In independent work, Chanda and Dalgaard (2008) and Vollrath (2009) perform a similar set a counterfactual exercises using capital stock data for agriculture and non-agriculture, and conclude that around 80% of international TFP differences can be accounted for by TFP differences in agriculture relative to non-agriculture.

While informative, these accounting exercises do not help us understand attempt *why* the agriculture sector exhibits so much more variation in productivity across countries than the non-agriculture sector, as our paper does. Another paper that endogenizes agriculture productivity differences is by Restuccia, Yang and Zhu (2008), who ask whether government-imposed barriers to the adoption of intermediate inputs, such as fertilizers, keep farm productivity low in the developing world. They find that in a plausibly calibrated model, barriers to intermediate adoption could explain as much as 25% of the agriculture productivity gap between the richest and poorest countries.

Gollin, Parente and Rogerson (2004) argue that barriers of a different sort, namely to capital accumulation, could be at the heart of low measured agricultural productivity in poor countries. They argue that these barriers could encourage workers to move from market production to home production, which is easier in the (rural) agriculture sector than (largely urban) non-agriculture sector. A central implication of their study is that much of agriculture production goes unmeasured, and that measured agriculture productivity differences overstate true productivity differences. Just how much agriculture output goes unmeasured is still very much an open question.

A final related paper is by Graham and Temple (2006), who argue that agriculture productivity differences across could stem from poor countries being in a different equilibrium than rich countries, where agriculture production is characterized by decreasing returns to scale and non-agriculture subject to increasing returns. Their hypothesis is that poor countries, with most workers in agriculture, have low agriculture productivity as a result since agriculture work is subject to decreasing returns to scale. Nevertheless, because of subsistence food needs, this constitutes an equilibrium. Rich countries, on the other hand, have few workers in agriculture, high agricultural productivity, and high overall income. They argue that, quantitatively, this theory can account for perhaps 15% to 25% of cross-country aggregate productivity differences.

3 Model of Agricultural and Non-Agricultural Productivity

In this section we develop a model of productivity in agriculture and non-agriculture relative to the aggregate. We show that in the model, sector-neutral efficiency differences across countries can generate relatively larger productivity differences in agriculture and relatively smaller differences in the non-agriculture sector than the efficiency differences themselves.

3.1 Households

There are measure one of agents, indexed by i, who differ by ability, as will be explained below. Preferences are given by

$$U^i = \log(c_a^i - \bar{a}) + \nu \log(c_n^i), \tag{1}$$

where c_a^i is food consumption, c_n^i is non-food consumption, $\bar{a}$ is a parameter representing a subsistence food requirement, and ν governs the relative taste for non-food consumption.

Each agent is endowed with one unit of time which she supplies inelastically to the labor market. Each agent is also endowed with a vector of abilities $\{z_a^i, z_n^i\}$ which represent the efficiency of one unit of labor in sectors a and n. The population density of abilities is drawn from a distribution $G(z_a, z_n)$ with support on the positive reals and positive variance for each ability. Agents earn wage income w^i, which is described in more detail below. The budget constraint is

$$p_a c_a^i + c_n^i \leq w^i \tag{2}$$

where p_a is the relative price of food, and the non-agricultural good is taken as the numeraire.

3.2 Production

There is a competitive market in each of the two sectors, and each has its own sector aggregate production function. Both sector technologies are freely available and operated by competitive entrepreneurs. The technologies are given by:

$$Y_a = A\tilde{L}_a \quad \text{and} \quad Y_n = A\tilde{L}_n \tag{3}$$

in agriculture and non-agriculture, where A captures sector-neutral efficiency, and $\tilde{L}_a$ and $\tilde{L}_n$ represent the total number of effective labor units employed in the two sectors. Let Ω^a and Ω^n denote the sets of agents electing to work in agriculture and non-agriculture. The sector aggregate labor inputs $\tilde{L}_a$ and $\tilde{L}_n$ are defined as

$$\tilde{L}_a \equiv \int_{i\in\Omega^a} z_a^i \, dGi \quad \text{and} \quad \tilde{L}_n \equiv \int_{i\in\Omega^n} z_n^i \, dGi$$

and represent the sum of all ability working in the respective sectors. Notice that our labor input differs from those of standard macro models in that ours sums up worker productivities, rather than workers themselves. The total number of workers in each sector is defined as

$$L_a \equiv \int_{i\in\Omega^a} dGi \quad \text{and} \quad L_n \equiv \int_{i\in\Omega^n} dGi.$$

3.3 Optimization and Equilibrium

Agents take as given prices and a wage schedule which maps abilities into sector-specific wage offers. The problem for an agent is first to pick which sector to work in, and then to maximize (1) subject to (2). Because of competition in production markets, the schedule of wages offered to a worker with abilities z_a^i and z_n^i is equal to:

$$w_a^i = p_a A z_a^i \quad \text{and} \quad w_n^i = A z_n^i \tag{4}$$

in the agricultural and non-agricultural sectors. A simple cutoff rule in *relative* ability determines the optimal occupational choice for each agent. Working in non-agriculture is optimal for agent i if and only if

$$\frac{z_n^i}{z_a^i} \geq p_a. \tag{5}$$

Thus, the agents that enter non-agriculture are those whose ability there is sufficiently high relative to their ability in agriculture. Let the resulting wage under the optimal sector choice be defined as $w^i \equiv \max\{w_a^i, w_n^i\}$.

The remainder of the agent's problem is standard, and optimal demands are:

$$c_a^i = \frac{w^i + \bar{a} p_a \nu}{p_a(1+\nu)} \quad \text{and} \quad c_n^i = \frac{\nu(w^i - \bar{a} p_a)}{1+\nu}. \tag{6}$$

Due to the subsistence consumption constraints, agents consume relatively more food when their wage is lower. The lower is ν, the higher the ratio of food to non-food consumption when the agent's wage is low.

An equilibrium of the economy consists of a relative food price, p_a, and allocations for all agents such that labor and output markets clear. Labor productivity in equilibrium is given by Y_a/L_a in agriculture, and Y_n/L_n, and represent the physical quantity of output produced per worker in each sector.

3.4 Relative Price of Agriculture Higher in Poorer Economies

In this section we show that, in equilibrium, the relative price of agriculture declines in the efficiency level, A.

Proposition 1 *Consider two economies, rich and poor, with efficiency terms A^R and A^P such that $A^R > A^P$. Then the relative price of agriculture is higher in the poor economy: $p_a^P > p_a^R$.*

To see the intuition for why p_a^P has to be higher than p_a^R, imagine in contradiction that they were the same. For expositional purposes, assume markets clear in the rich country. Then, by (5), the sector labor supply cutoffs would be the same in both countries, and hence so would the share of workers electing to supply labor in the agriculture sector. But because of the subsistence food requirement, the poorer economy demands a much larger fraction of food. Hence output markets would not clear in the poor economy. In order to induce enough workers to supply labor in agriculture in the poor economy, it must be true that p_a^P is greater than p_a^R.

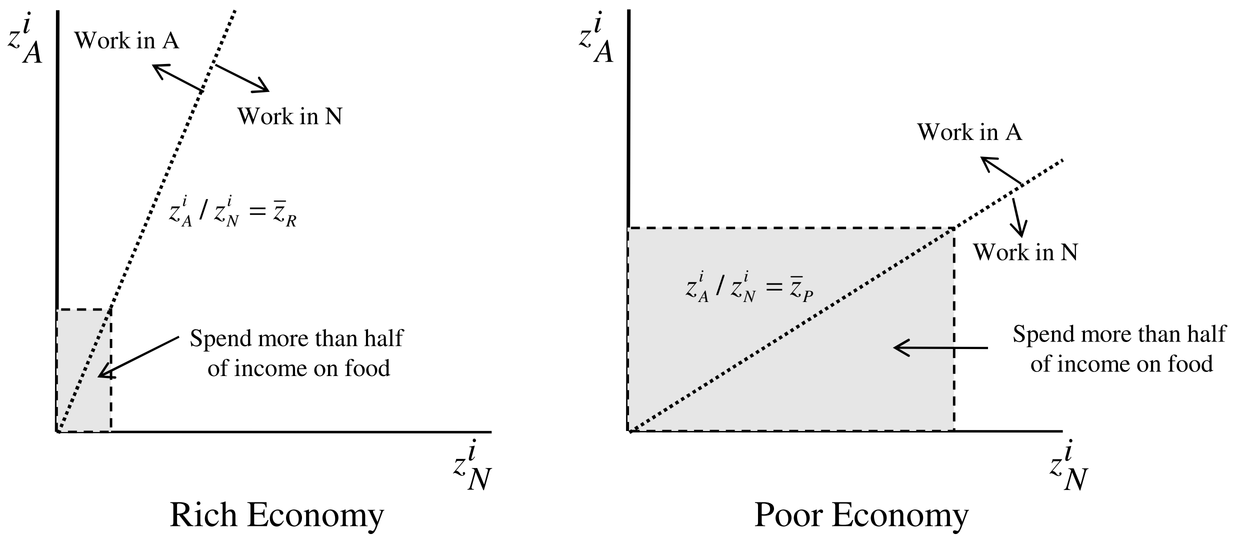

Figure 1: Optimal Sector Choice in Rich and Poor Economies

Figure 1 illustrates optimal sector choice in equilibrium. Each point on the figure represents one conceivable draw of (z_a, z_n), corresponding to a pair of sector-specific abilities. The dotted lines stemming from the origin describe the set of ability pairs for which agents are indifferent between the two sectors, i.e. when z_n^i/z_a^i equals p_a^P and p_a^R respectively. Points above the lines represent agents for which working in sector n is optimal, and points below the lines meaning that working in a is optimal. As in Proposition 1, because $p_a^P > p_a^R$, more agents work in non-agriculture in the richer economy. The shaded regions describe the set of agents that spend more than half their income on food.[1] The poor economy has a larger fraction of such agents because of the subsistence food requirement.

3.5 Illustrative Example: Two Types

In this section we illustrate, using a simple two-type example, the intuition for how productivity differences can be larger in agriculture and smaller in non-agriculture given sector-neutral

[1]The choice of one half income spent on food is arbitrary, and just meant to convey the higher food share in the poor country.

productivity differences. Let the two types be called *strong* and *smart*, and be endowed with ability vectors $\{\sigma, 1\}$ and $\{1, \sigma\}$, respectively, where $\sigma > 1$ governs the extent of comparative advantage in one sector versus the other. Put simply, the strong types have a comparative advantage at farming, and the smart types have a comparative advantage at non-agricultural work. Let half the population be of the strong, and the other half smart.

3.5.1 Productivity Differences Larger in Agriculture than Non-Agriculture

One can readily see that when efficiency is sufficiently low, say at a level A^P, then all the strong types and at least some of the smart types work in agriculture, and only smart types work in non-agriculture. The smart types that work in agriculture are induced to do so by being offered a wage equal to their marginal product in non-agriculture, namely σ. Thus, the price of agriculture goods must be $p_a^P = \sigma$. Letting α be the measure of smart agents working in agriculture, sector productivities are given by

$$Y_a^P/L_a^P = A^P \frac{(1/2)\sigma + \alpha}{1/2 + \alpha} \quad \text{and} \quad Y_n^P/L_n^P = A^P \sigma. \tag{7}$$

Next, note that for efficiency sufficiently high, only strong types are left in agriculture, whereas some strong and all smart types work in non-agriculture. In this case the strong types are induced to work in non-agriculture only when the price of agriculture is $p_a^P = 1/\sigma$. In this case, strong types working in agriculture produce σ units at price $1/\sigma$ for a value of 1, and strong types working in non-agriculture produce 1 unit of the non-agriculture good at (normalized) price 1. Now, letting β be the measure of strong types working in non-agriculture, productivities are given by

$$Y_a^R/L_a^R = A^R \sigma \quad \text{and} \quad Y_n^P/L_n^P = A^R \frac{\beta + (1/2)\sigma}{\beta + 1/2}. \tag{8}$$

We can now see that productivity differences are larger in agriculture than the exogenous efficiency differences. Comparing (7) and (8), we have

$$\frac{Y_a^R/L_a^R}{Y_a^P/L_a^P} = \frac{A^R}{A^P} \frac{(1/2)\sigma + \sigma\alpha}{(1/2)\sigma + \alpha} > \frac{A^R}{A^P}. \tag{9}$$

Furthermore, we can see that non-agriculture productivity differences are smaller then the efficiency differences:

$$\frac{Y_n^R/L_n^R}{Y_n^P/L_n^P} = \frac{A^R}{A^P} \frac{\beta + (1/2)\sigma}{\sigma\beta + (1/2)\sigma} < \frac{A^R}{A^P}. \tag{10}$$

The intuition for these results is that the two economies have different degrees of specialization, with lower average ability farm workers in the poor economy, and lower average ability non-agriculture workers in the rich economy. These specialization differences are induced by the sector-neutral efficiency differences and the subsistence food constraint, which induce those less able in agriculture to work in agriculture when efficiency is low.

3.5.2 ...and the Productivity Differences are Increasing in Ability Differences

One important feature of this illustrative example is that the size of the sector productivity differences are increasing in the extent of ability differences. If both types were equally able in each sector, i.e. $\sigma = 1$, then by (9) and (10) we can see that sector productivity ratios are equal to exactly $\frac{A^R}{A^P}$, which are the aggregate gaps. On the other hand, as σ increases, the sector differences become even larger in agriculture, and even smaller in non-agriculture. This simple intuition will be of key importance in the quantitative section to follow. The larger the parameterized model's ability differences across agents, the larger the model's predictive power.

3.5.3 ...but Only When Ability Differences Give Rise to Comparative Advantage

We now illustrate that one requirement for productivity differences to be larger in agriculture than the aggregate is that there be comparative advantages in ability across agents. In other words, it must be true that some agents have a comparative advantage in one sector over other agents. Without comparative advantage differences, individuals potentially differ only in absolute ability, and our mechanism is shut down.

To see this, consider a variant of the illustrative model above where agents come in two different types, namely *good* and *bad*. Let their ability vectors be $\{1, \sigma\}$ and $\{\gamma, \gamma\sigma\}$ where $\gamma > 1$. In this case there is no sense of comparative advantage, since the good type is simply a factor γ times as productive in each of the two sectors.

Under this assumption on ability, one can see that the wage offers to bad and good agents must be $A\sigma$ and $A\gamma\sigma$, respectively, for given efficiency A. The price of the agriculture good is $p_a = \sigma$ for any A, and more importantly, good and bad agents are each indifferent between working in the two sectors. Thus, the composition of good and bad agents in each sector is indeterminate (although the total number of effective labor units in each sector is pinned down for any A) and therefore it need not be true that agricultural productivity differences are larger than A differences across economies with different A values. This simple intuition will also carry through to the quantitative section to follow, when we parameterize the extent to which sector ability draws are correlated across individuals. The higher the correlation, the lower the predictive power of the model.

4 Quantitative Analysis of the Model

In this section we parameterize the model and then asses its quantitative importance for understanding agriculture and non-agriculture labor productivity differences across countries. We also assess the model's predictions for cross-country and historical U.S. data on shares of labor and GDP in agriculture, relative prices of agricultural goods, plus the current cross-section of wages in the U.S.

To parameterize the model we must select a distribution of ability, $G(z_a, z_n)$, plus values for the two taste parameters $\bar{a}$ and θ; the efficiency term A can be normalized to 1 for the United States. We discuss each in turn.

4.1 Parameterization of Ability Distribution

We set the joint distribution of abilities to be:

$$G(z_a, z_n) = \exp\left\{ -\left[\left(-\log F(z_a)\right)^{\rho} + \left(-\log H(z_n)\right)^{\rho} \right]^{\frac{1}{\rho}} \right\}$$

$$\text{with } \; F(z_a) = e^{-z_a^{-\theta_a}} \text{ and } \;\; H(z_n) = e^{-z_n^{-\theta_n}}.$$

The functions $F(z_a)$ and $H(z_n)$ are the CDFs of Fréchet random variables.[2] The lower are θ_a and θ_n, the higher is the variance of ability in agriculture and non-agriculture. Dependence of the ability draws is induced using the function $\exp\left\{ -\left[\left(-\log u\right)^{\rho} + \left(-\log v\right)^{\rho} \right]^{\frac{1}{\rho}} \right\}$, which is known as a Gumbel copula.[3] The parameter $\rho \in [1, \infty)$ determines the extent of dependence, with a higher ρ representing more dependence between draws.[4]

Two natural questions are: why we chose a parametric form for the ability distribution, as opposed to something nonparametric, and why we chose this particular parametric form. The answer to the first question is that the Roy model, in spite of its apparent simplicity, cannot be identified from cross-sectional wage data without making assumptions on the functional form of the ability distribution (Heckman and Honore, 1990.) Because one only observes the

[2] This distribution has been used to parameterize Ricardian models of international trade originating with Eaton and Kortum (2002). It is also very convenient. In trade models, the Fréchet distribution yields a log-linear gravity equation relating trade flows to structural parameters. Similarly in our framework with a common θ parameter, our model yields a log-linear equation relating employment shares to the relative price of agriculture goods with elasticity $\rho\theta$.

[3] If $\theta_a = \theta_n$, then this distribution is the same as the "multivariate" Fréchet used in Ramondo and Rodríguez-Clare (2009) and discussed in footnote 14 of Eaton and Kortum (2002). However, unlike in these contexts, we show how to identify the dependence parameter.

[4] Copulas can be used to create multivariate distributions out of arbitrary univariate distributions. See e.g. Cherubini, Luciano and Vecchiato (2004).

maximum of each agents' draws, but not both draws themselves, if ability distributions are allowed to take on an arbitrary form, there are many distributions which can generate a given set of observations on wages and sector choices by individuals.

There are three main justifications for this particular parametric form of the ability distribution. First, it parsimoniously allows for (potentially different) ability dispersion in each sector as well as dependence between ability draws. As will be shown in the following section, the three parameters of the distribution (θ_a, θ_n and ρ) can each be disciplined in a transparent way using a single cross section of wages for agricultural and non-agricultural workers.

Second, the choice of Fréchet distributions for ability in each sector contains a sensible economic interpretation, which is as follows. The Fréchet distribution is an extreme value distribution, representing the distribution of the maximum of independent draws from some underlying distribution. Thus the draw z_n^i, for example, can be thought of as the maximum of household i's ability draws in a large set of distinct non-agricultural tasks. A similar interpretation can be given to z_a^i.[5]

Finally, Fréchet distributions for ability yield wage distributions – for the economy as a whole and by sector – which closely resemble their empirical counterparts, as we demonstrate in the following section. In particular, the model delivers wage distribution tails that mimic the data extremely closely, a dimension along which other distributions fail. Heckman and Sedlacek (1985) and Heckman and Honore (1990), for example, argue that a Roy model with lognormal ability distributions generates tails which are too thin compared to the data. Since our results are sensitive to the size of the tails of the ability distribution, it is essential that tails in the model are in line with the data.

To calibrate the ability distribution parameters, our basic strategy is to use cross sectional wage data from the United States. Formally, we jointly determine θ_a, θ_n and ρ to match three moments, the standard deviations of log wages in agriculture and non-agriculture plus the ratio of average wages in agriculture and non-agriculture. Intuitively, the our calibration strategy is as follows. The θ_a and θ_n terms determine the variation in ability across individuals, with higher θ's resulting in lower variation in abilities. Because wages are set equal to the value of marginal products, variation in ability maps into variation in wages across agents. Thus, observed wage variation in agriculture and non-agriculture wages are key in disciplining the parameters θ_a and θ_n.

Next, ρ is pinned down by the average wage in agriculture relative to non-agriculture, with a lower relative agriculture wage implying a higher ρ. We argue this point in Figure 2(a). Figure

[5]By the extreme value theorem, the maximum of a sample of i.i.d. draws from any distribution converges in distribution to one of three extreme value distributions: the Fréchet, the Gumbel, or the Weibull. See e.g. Kotz and Nadarajah (2000).

2(a) plots a simulation with each point in the plane depicting the two wages that a household can receive. If a household is to the left of the 45^o line than it is optimal to work in agriculture and vice versa. In both figures, only five percent of the workers chose agriculture. Also plotted is the log mean wage of workers in agriculture and non-agriculture across the two scenarios. Figure 2(b) plots a simulation with no correlation. The key thing to notice is that workers

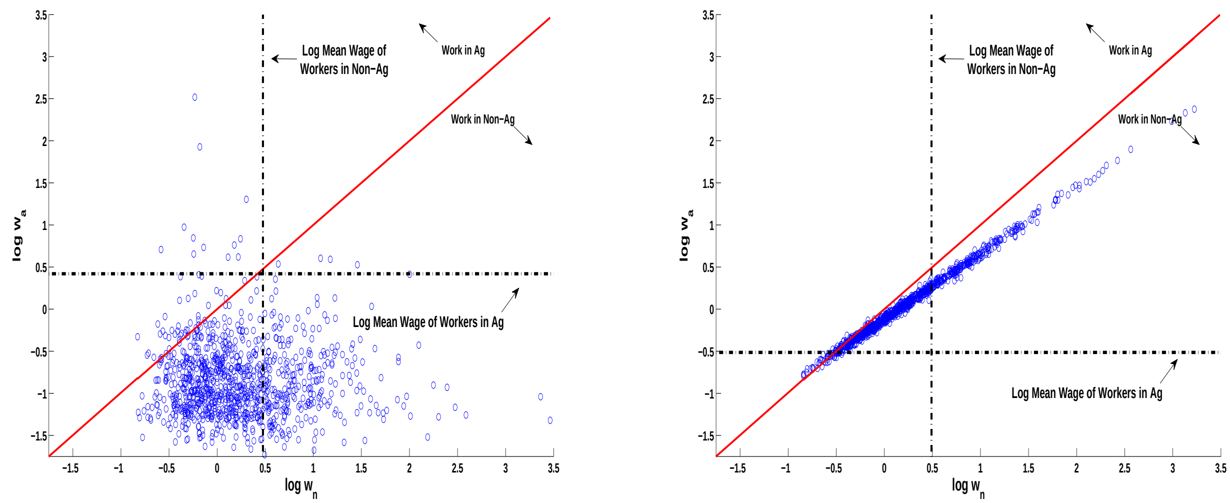

(a) No Dependence, $\rho = 1$.

(b) Positive Dependence, $\rho = 20$.

Figure 2: Wages Offers From Model with Dependence and No Dependence

in agriculture earn on average lower wages (compared to non-agriculture) in the model with dependence relative the model with no dependence. Mechanically this can be seen how the horizontal dashed line (indicating the log mean wage in agriculture) shifts down across the two scenarios, while the vertical dashed line (indicating the log mean wage in non-agriculture) stays essentially the same across the two scenarios. What this implies is that the dependence parameter ρ affects the ratio of the average wages in the two sectors. Hence we pick ρ to target this moment.

We get our cross-sectional data from the U.S. Current Population Survey (CPS) for 2007. Following the study of U.S. wage inequality by Heathcote, Perri and Violante (2009), we take all individuals between 25 and 60 who have non-missing data on income and hours worked. We restrict the sample further to include only workers averaging at least 35 hours per week of work, and only those earning at least the Federal minimum wage. These restrictions provide more conservative estimates of cross-sectional wage variance, which will lead to more conservative variances of ability in the parameterized model (and hence less predictive power for our mechanism). We calculate that the standard deviation of log wages for agriculture workers is 0.46, while in non-agriculture the standard deviation of log wages is higher at 0.57. The average wage in agriculture is 0.77 times the average wage in non-agriculture.

These figures imply (along with the preference parameters calibrated below) choices of $\theta_a = 2.88$, $\theta_n = 2.25$ and $\rho = 1.03$. Since ρ itself is hard to interpret (it runs from 1 to ∞), we computed the Kendall rank correlation coefficient to be 0.035. This suggests that there is a moderate amount of correlation across individuals in the their abilities across the two activities. The estimates of θ_a and θ_n mean that there is more variance in ability in non-agriculture work than in agricultural work, again reasonable given that non-agriculture work encompasses so many more types of tasks.

4.2 Parameterization of Preferences

For the preference parameters, we pick ν to match average share of labor in agriculture across countries in the top 10th percentile of the income distribution. The resulting parameter implies a a long-run food expenditure share of 2.7%. Caselli and Coleman (2001) and Duarte and Restuccia (2009) pick a value of 1%, while Restuccia, Yang and Zhu (2008) pick a value of 0.5%. Admittedly, our model's results are sensitive to the choice of this value, with lower values allowing us to explain more of the variation in agriculture and non-agriculture labor productivity. For this reason we stick with a more conservative value relative to others in the literature.

We set $\bar{a}$ to match a subsistence consumption need of 34% of average income in a model country with 7.5% of the U.S.'s per capita GDP. This is consistent with the independent estimates of subsistence food consumption requirements of Rosenzweig and Wolpin (1993), and Atkeson and Ogaki (1996), both of which use panel data from a sample of rural households from India (which had 7.5% of the U.S. per capita GDP in 1984).[6]

4.3 Quantitative Predictions for Sector Productivity Differences

With the model parameterized, we now ask what it predicts quantitatively for agriculture productivity differences and non-agriculture differences in the cross section of countries. Specifically, we solve the model over a range of A values covering the world income distribution, and compute its predictions for relative output per worker in the aggregate and the two sectors.

Table 2 shows the model's predictions for the ratio of the 90th to 10th percentile of countries in the model and data. The differences in aggregate output per worker (expressed as GDP per worker at Gheary-Khamis international prices) is a factor of 22 in the model and data by

[6] Rosenzweig and Wolpin estimate a subsistence requirement of 1,469 rupees per agent per year. Townsend (1994) reports that average agent size in the sample is 6.7 and that average income per person in the Indian sample is 635 rupees.

Table 2: Labor Productivity Differences

Sector	Ratio of 90th-10th Percentile		
	Data	Model	Percent Explained
Agriculture	45	40	78
Aggregate	22	22	-
Non-Agriculture	4	10	67

Data Source: Caselli (2005)

construction. If the model had no predictive power, it would predict that agricultural and non-agricultural productivity gaps would be equal to the aggregate gaps. Instead, the model predicts agriculture output per worker differences should be a factor of 40, and in non-agriculture it predicts a factor of 10 difference. In the data these ratios are (as described in Section 2) a factor of 45 and 4 respectively. The third column of the table shows that this corresponds to the model explaining 78% the difference between the agricultural productivity gap and the aggregate gap, and 67% of difference between the non-agricultural gap and the aggregate. The results in Table 2 show that differences in patterns of specialization along with sector-neutral productivity differences are quantitatively important to understanding why agriculture (non-agriculture) sector productivity differences are so much larger (smaller) than aggregate differences.

Table 3 illustrates the model's predictions for developing countries with relatively higher average income. Specifically, it shows the model's prediction for the 90th-50th ratio and 90th-25th ratio. In the latter case aggregate productivity in the model and data differ by a factor of 9.4, again by construction. In the 90th-25th case, the model predicts a factor of 16.5 in agriculture and 7.6 in non-agriculture, compared to 31.1 and 2.7 in the data. The model predicts 33% and 27% of the agricultural and non-agricultural productivity differences, relative to the aggregate, as in the data, which is still large, but substantially lower than the 90th-10th percentile ratio.

In the 90th-50th case, the model fares much worse. The aggregate differences are chosen to be a factor of 3.1 as in the data. The model predicts differences in agriculture and non-agriculture of 3.9 and 3.0, compared to 11.1 and 1.9 in the data. This amounts to explaining just 10% and 8% of the sector differences compared to aggregate differences, respectively.

Why does the model fare so much less successfully when explaining differences between rich countries and those at intermediate income levels? The answer has to do with differences in the shares of labor in agriculture. Consider the case of the 90-50 differential. The percent of workers in agriculture in the 50th percentile country is just 9%, compared to 2% in the 90th percentile country. Thus, the average productivity of workers in agriculture is only slightly lower in the 50th percentile country. In contrast, in the 10th percentile country, 74% are in

Table 3: Labor Productivity Differences – Intermediate Income Levels

Sector	Ratio of 90th.-25th Percentile		Percent Explained
	Data	Model	
Agriculture	31.1	16.5	33
Aggregate	9.4	9.4	-
Non-Agriculture	2.7	7.6	27
	Ratio of 90th-50th Percentile		
	Data	Model	
Agriculture	11.1	3.9	10
Aggregate	3.1	3.1	-
Non-Agriculture	1.9	3.0	8

Data Source: Caselli (2005)

agriculture, and thus the average worker has substantially lower productivity than the average agricultural worker in the 90th percentile country. Hence, the model's explanatory power is larger for differences between the the richest countries and the poorest countries than for the richest and middle income countries.

4.4 Assessing The Model's Other Quantitative Cross-Country Predictions

The model generates large productivity differences in agriculture relative to non-agriculture across countries, at least between the richest and poorest countries. We now ask whether it is successful in matching other relevant data. In particular, we compute the model's cross-country predictions for the share of labor in agriculture, the share of GDP spent in agriculture, and the relative price of agriculture and assess whether they are quantitatively consistent with the data.

Figure 3(a) shows the model's predictions for the cross-section of countries for the share of labor in agriculture, along with the actual data. The horizontal axis displays purchasing-power parity GDP per worker for 2000 (on a log scale), and the y-axis displays the percent of workers in agriculture. As in the data, our model predicts that the poorest countries should have shares in the range of 70% to 90% of all workers, down to less than 10% for the richest. The model also captures the convex nature of this curve, which is driven in the model by the concavity of preferences along with the subsistence constraint in food. We conclude that this feature of the data is successfully captured by our model.

Next we turn, in Figure 3(b), to the model's predictions for the share of GDP in agriculture.

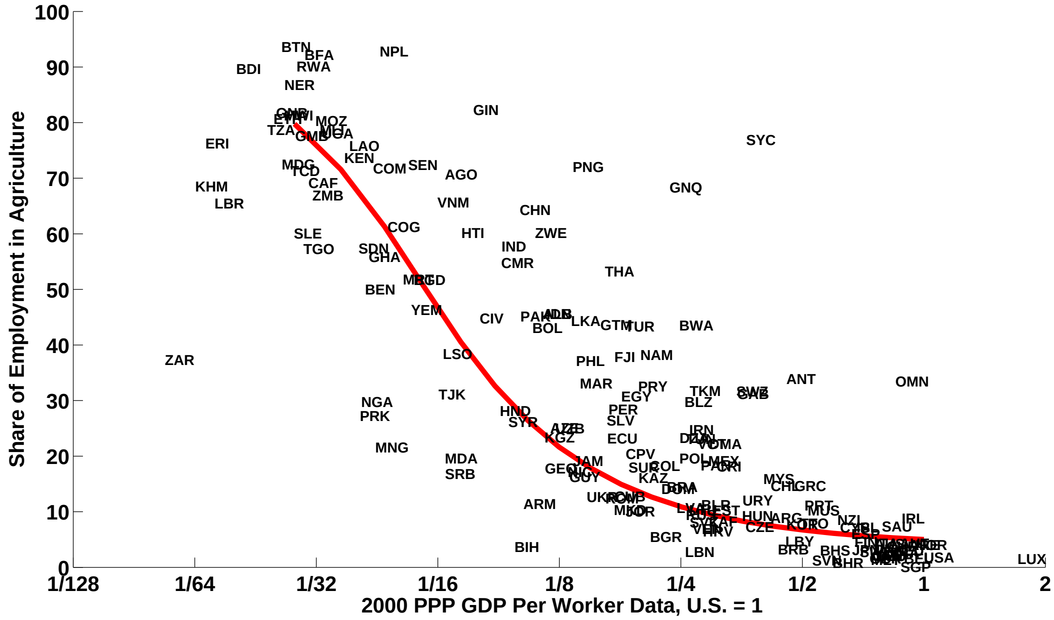

(a) Employment Share in Agriculture, Data and Model

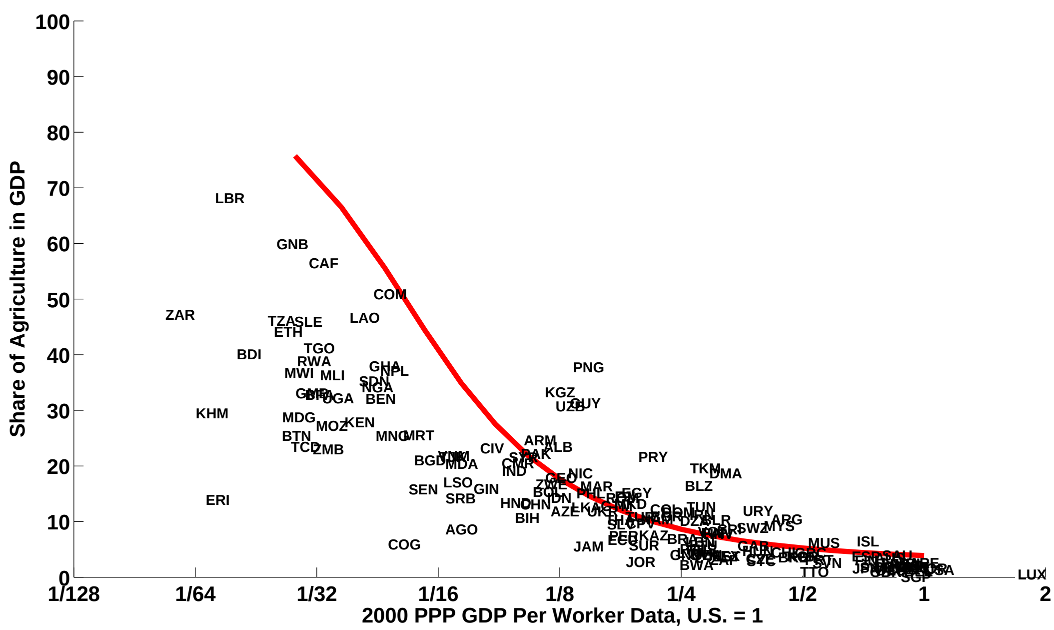

(b) GDP Share in Agriculture, Data and Model

Figure 3: Employment and Agriculture Shares, Data and Model

While similar to the labor shares shown above, note that in the data the GDP shares in agriculture are systematically lower than the labor shares in agriculture. In Kenya, for example, agriculture employs 74% of the workers but produces just 28% of GDP. While our model does a reasonable job of capturing GDP shares for agriculture in the countries with around 1/8 the U.S. income level or higher, it substantially over-predicts the GDP share in agriculture in countries with lower income.

One reason for the model's inconsistencies with the data in this dimension may be that agricultural GDP itself is mis-measured in the poorest countries. If households spend much of their time in home production of agricultural goods, then measured GDP of agriculture will understate true agricultural output (Gollin, Parente, Rogerson, 2004).

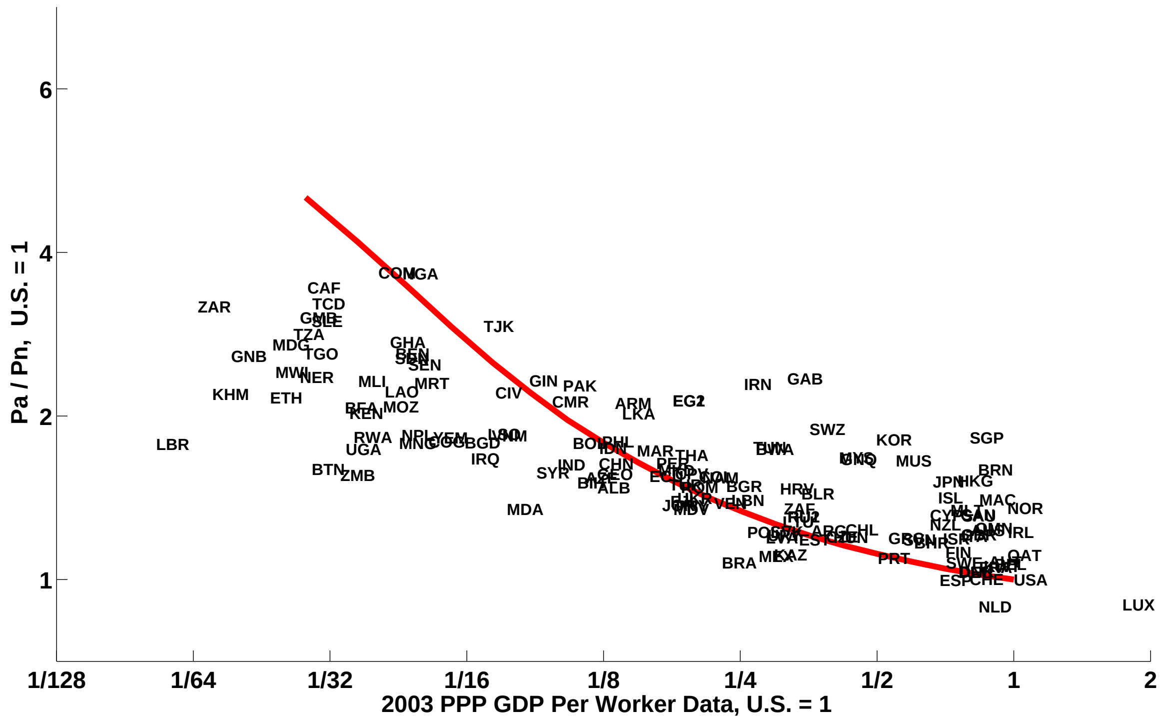

Figure 4: Relative Agriculture Prices in Model and Data

An important prediction of our model is that the relative price of food is higher in poor countries than rich countries. We now ask whether this prediction is borne out in the data, and whether the model is quantitatively consistent with the cross-country relationship between relative food prices and income per capita. Figure 4 plots this relationship. The vertical axis contains the relative price of agricultural goods (expressed in log base 2) with the U.S. value normalized to one and the horizontal axis plots real GDP per worker relative to the U.S. also in log base 2 scale. Our data on relative food prices are constructed using 2005 data available from the International Comparison Programme (see Appendix B for details).

As can be seen in the Figure, relative food prices are indeed higher in lower income countries, as predicted by Proposition 1. The ratio of relative food prices in the 10th percentile of the country income distribution to 90th percentile is about 2.5. In the model it is around 4. We conclude that while model performs well except for the poorest countries, where it over-predicts the relationship found in this data.

One possible concern with this ICP relative price data is that since they are consumer prices, not producer prices, they could reflect differences in distribution services across countries. In particular, it could be the case that the higher relative price of food in poor countries is largely due to higher distribution costs for food in those countries, not less efficient production in agriculture production relative to non-agriculture, as in the model.

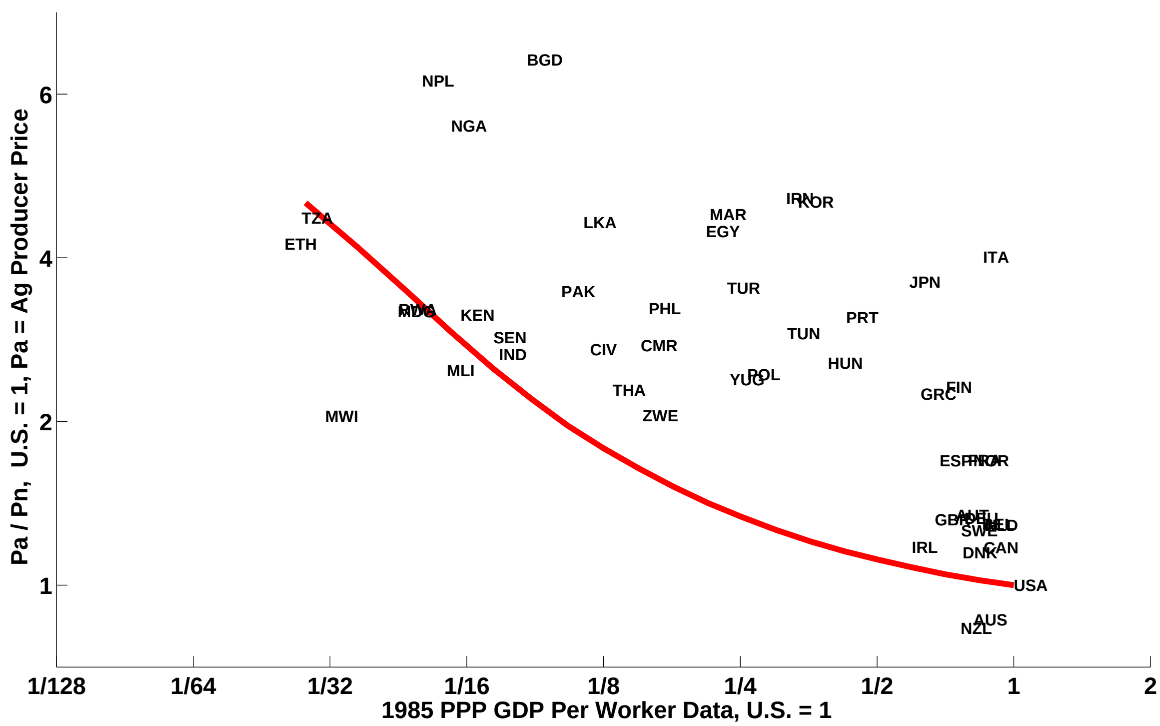

Figure 5: Relative Agriculture Producer Prices in Model and Data

To address this concern we computed the relative price of agriculture for all countries for which producer prices of agriculture goods are available. Our data source is the 1985 FAO food producer price data, explored in detail by Adamopoulos (2008), and used by Caselli (2005) to construct sector productivity measures. For the prices of non-agriculture goods we use the consumer price data for the corresponding countries available in the 1985 Penn World Tables. We end up with 60 countries with reasonably broad variance in per capita income.

Our results using producer prices of agriculture are below. According to Figure 5, one can see that relative prices of food appear even higher once producer prices are used. This is consistent with the finding of Adamopoulos (2008) that distribution margins for food are moderately higher in richer countries than poor countries.

4.5 Predictions for U.S. Cross-Sectional Distribution of Wages

A final implication of our model worth testing is the entire distribution of wages within an economy. In our calibration we made parametric assumptions regarding the distribution of abilities and calibrated the parameters to target several moments of the wage distribution. One may be concerned that our functional form choice has unreasonable implications for relative to the entire distribution of wages. Figure 6(a) plots the empirical cumulative distribution function from the U.S wage data and from data generated by our model. They track each other very closely. Figure 6(b) and 6(c) plot the same relationship for only those workers in agriculture and non-agriculture. The model performs reasonably well, particularly in its upper tails, which closely mimic those in the data. The only substantiative deviation is that our calibration suggests there are slightly more low wage workers in than seen in the data. Overall, the Figures show that our parametric assumptions on the ability distribution yield realistic wage distributions.

4.6 More Conservative Calibration

While wage variation in the model arises only because of variation in efficiency of labor across individuals, some economists argue that wage variation in the data is due in part to factors unrelated to productivity, such as market imperfections resulting from search frictions or transitory variation in wages. The largest estimate of the importance of these imperfections (that we could find) is that of Postel-Vinay and Robin (2002), who estimate that around one-half of the variance in log wages is due to market imperfections. Alternatively, Guvenen (2009) decomposes wage variation into "permanent" and "idiosyncratic" components. He finds that the idiosyncratic component accounts for up to 45 percent of overall variance in log wages. In an effort to be as conservative as possible in determining the extent of ability differences across individuals, we follow the Postel-Vinay and Robin estimate, implying that the standard deviation of the log wages due to ability differences is 0.32 in agriculture and 0.40 in non-agriculture. We follow the our previous approach by calibrating ρ to best fit the average wage in agriculture relative to non-agriculture.

Calibrating the ability distribution parameters to these moments results in $\theta_a = 3.1$, $\theta_n = 4.4$, and ρ very close to zero (recall that it is not defined exactly at zero). This choice of ρ leads to a slight overestimate the relative average wage suggesting a need for negative dependence in

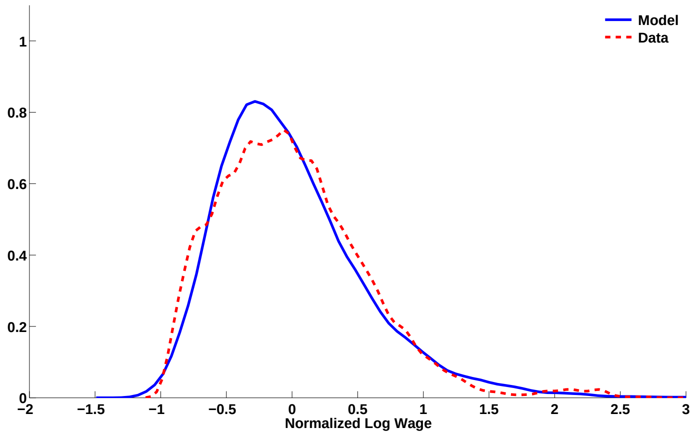

(a) Economy Wide Wage Distribution

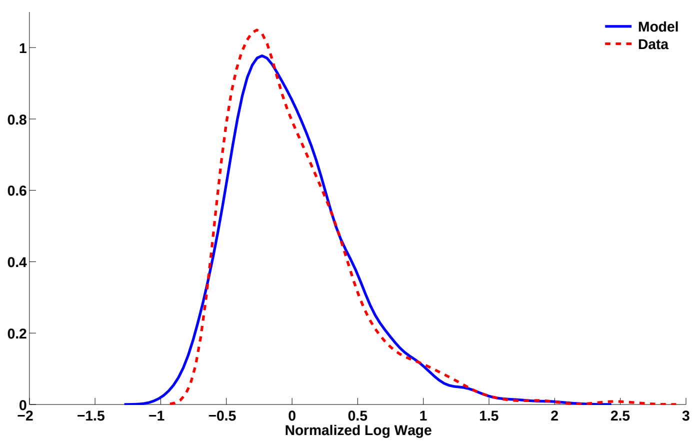

(b) Agriculture Worker Wage Distribution

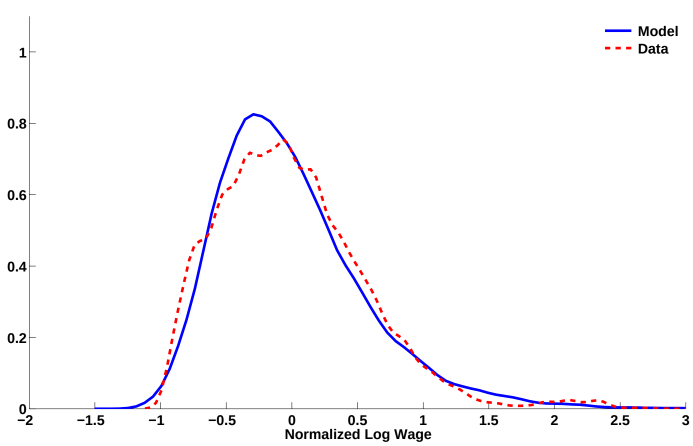

(c) Non-Agriculture Worker Wage Distribution

Figure 6: Wage Distributions, Model and Data

order to match the relative average wage. Since negative dependence increases the model's explanatory power, we leave ρ close to zero, and note that the predictions of this exercise will be even lower than they would have been with ρ chosen to match the relative wages exactly.

With this conservative calibration, we now ask what it predicts quantitatively for agriculture productivity differences and non-agriculture differences in the cross section of countries. Table 4 presents the results for countries in the 90th and 10th percentile of the income distribution.

Table 4: Labor Productivity Differences, Conservative Calibration

Sector	Ratio of 90th-10th Percentile		
	Data	Model	Percent Explained
Agriculture	45	30	35
Aggregate	22	22	-
Non-Agriculture	4	12	56

Data Source: Caselli (2005)

The model predicts agriculture output per worker differences should be a factor of 30, and in non-agriculture it predicts a factor of 12 difference. In the data these ratios are a factor of 45 and 4 respectively. The third column of the table shows that this corresponds to the model explaining 35% of agricultural differences and 56% of non-agricultural differences, relative to aggregate differences. The key result from this exercise is that our mechanism still explains around half of the variation in sectoral productivity gaps even when we feed into a model which generates only one half the observed variation in wages.

Similar to the baseline calibration, the model is not as successful at explaining gaps for intermediate income levels simply because differences in aggregate labor allocations are not as large. The other predictions of the model are quantitatively similar as well. The model correctly captures the employment share and agriculture share in GDP. The model's predictions for relative prices improve with the model generating and elasticity of relative prices with respect to income level of -0.23. In the data this elasticity is -0.21.

4.7 Testing the Model Using U.S. Historical Evidence

One natural testing ground for the model's quantitative predictions is the United States over the last 150 years, during which time the country experience sustained growth in income per capita and a structural transformation featuring a dramatic reduction in the importance of agricultural production. In this section we compare U.S. time series on income per capita, the share of GDP and employment in agriculture, and the relative price of agricultural goods to our model's predictions.

Figures 7(a) and 7(b) shows the share of employment and GDP in agriculture in the model and U.S. historical data. Each dot represents a year, and the x-axis and y-axis are identical in scale to those in the cross-country comparison earlier. As can be seen, the model is largely in line with the data in both cases, with higher shares of GDP and employment in agriculture when income per capita is lowest.

While the model is quite close to the data in terms of GDP share in agriculture, it under-predicts the share of labor in agriculture for much the time period. For the period 1880 through 1940, the actual share of workers in agriculture was between 20% and 40%. In the model, however, the model predicts between 15% and 30%. One possible way of reconciling the model's under-prediction is that the U.S. was a major exporter of food over this period, unlike in the (closed economy) model. For the post-war period, the model's predictions are much closer to the data, with shares in the range of 15% and below.

Turning to prices, Figure 8 shows that earlier in the United States history, when GDP per head was lower, the relative price of agricultural goods was higher. Relative agriculture prices were around twice as high around the turn of the century (when per capita GDP was in the range of $1/8$ of its current value) as they were in1990. All in all, the model's price predictions line up reasonably well, with moderate under-prediction of the price differences for most periods. We conclude from this section that, in spite of its simplicity, the model is largely consistent with the U.S. historical evidence.

5 Historical Evidence: Males versus Females

In this section we provide some direct evidence in support of our theory. While many dimensions of ability heterogeneity are not observable, along two particular observable dimensions, namely age and sex, there is concrete evidence of ability differences in agricultural and non-agricultural tasks. Historians and development economists have argued that women and children generally have a comparative disadvantage at farm work than adult men. As one piece of evidence, Goldin and Sokoloff (1982, 1984) show that wages were much lower for women in farm work in the United States, earning roughly one third to one half as much as men in farm work in the nineteenth century, with smaller wage differentials in manufacturing work. Foster and Rosenzweig (1996) provide complementary evidence from a sample of farmers from the Philippines, among which they estimate males to have an absolute productivity advantage in

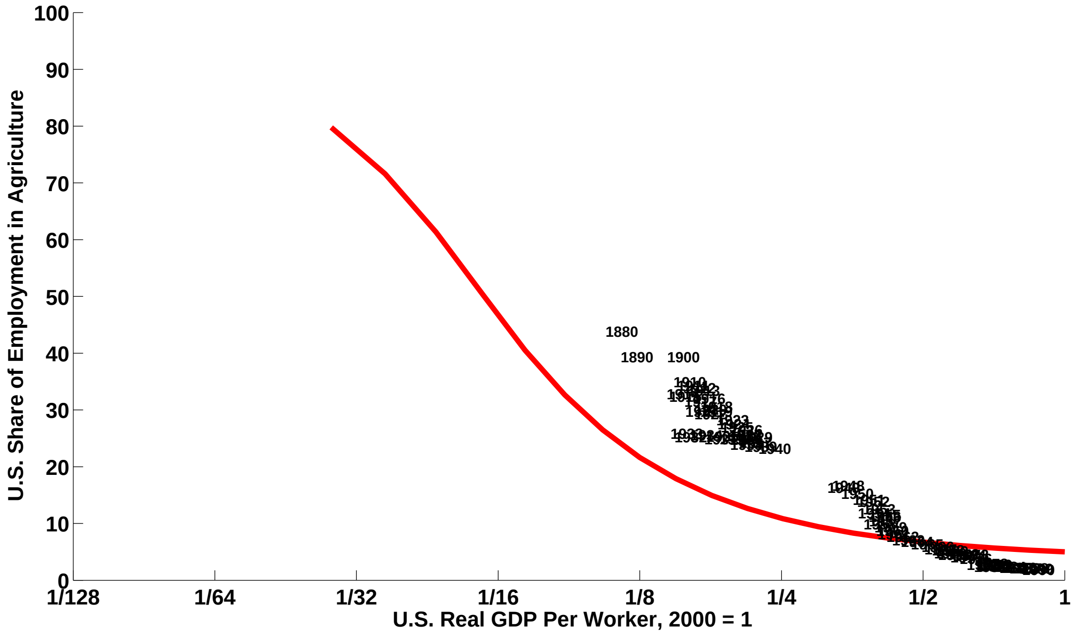

(a) U.S. Employment Share in Agriculture, and Model

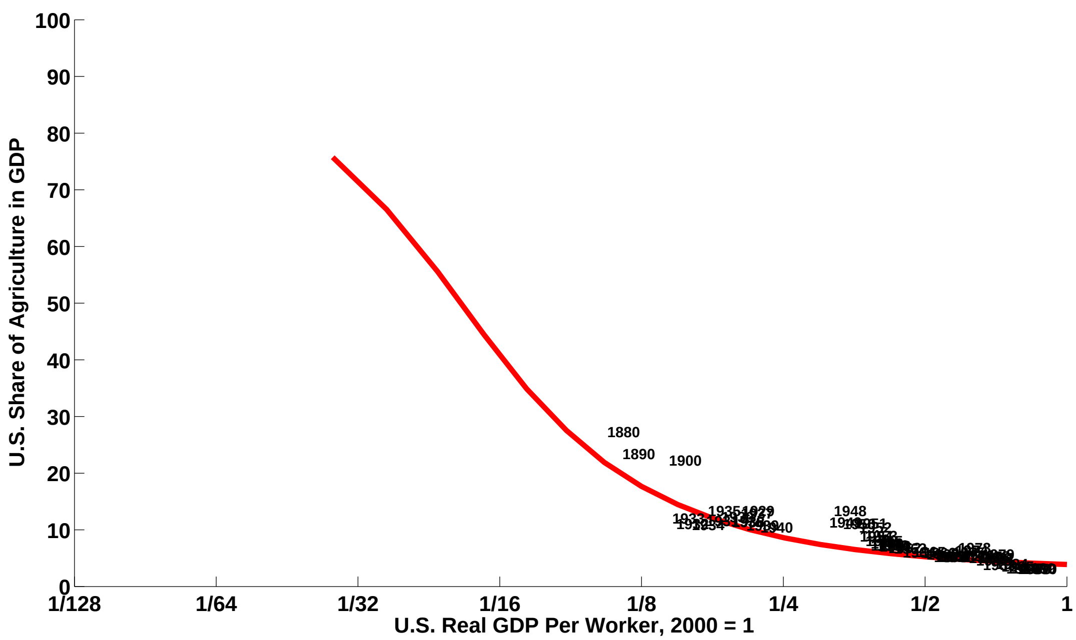

(b) U.S. GDP Share in Agriculture, and Model

Figure 7: U.S. Historical Agriculture Shares of Employment and GDP, and Model

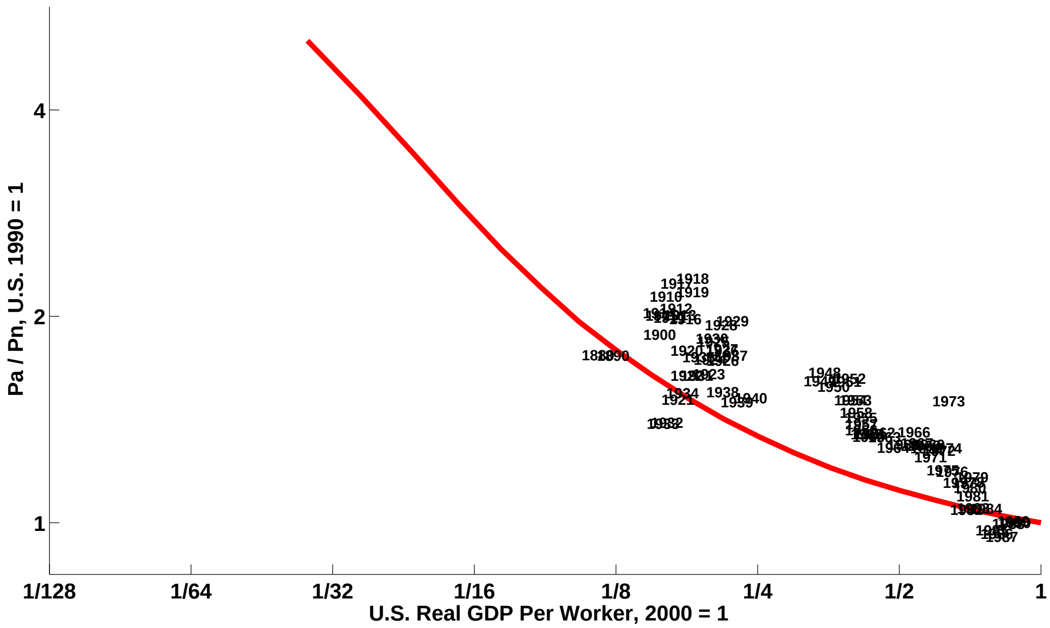

Figure 8: Relative Agriculture Producer Prices in U.S. and Model

several types of agricultural tasks.[7]

Our theory thus predicts that women and children should have been the first to move off the farms and into non-agricultural work. This is in fact what happened. In Britain, according to Allen (1994), the fraction of farmers that were women or children declined substantially during Britain's industrial revolution. Table 5 shows Allen's calculations for the composition of farm workers in England and Wales between 1700 and 1851. In 1700, a full 62.0% of farm workers in England were women and children, with the balance adult men. By 1800 this percent fell to 55.3%, and by 1851 it was down to 36.3%. For just women, the same figures are 32.5% in 1700, 30.3% in 1800, and 26.8% in 1851. Men, on the other hand, went from representing just over one third of farm workers to just under two-thirds.

Table 5: Composition of English Farm Workers

	1700	1800	1851
Men	38.3 %	44.7 %	63.7 %
Women	32.5 %	30.3 %	26.8 %
Women and Children	62.0 %	55.3 %	36.3 %

Data Source: Allen (1994)

[7] They estimate a one-factor model with two tasks: ploughing and weeding. They find that men are more productive at both, with a larger productivity difference in ploughing.

Goldin and Sokoloff (1982, 1984) provide evidence that this pattern held in the United States as well, using evidence from the manufacturing sector, which was a major component of the non-agricultural economy in the nineteenth century. In 1820, in the Northeast United States, roughly 55% of manufacturing workers were women and children. By 1890, this figure was down to 21%. The interpretation given by Goldin and Sokoloff is that as manufacturing work became available, women took manufacturing jobs at a faster rate than men, who stayed in agriculture work relatively longer.

Furthermore, Goldin and Sokoloff argue that the primary reason women moved into manufacturing relatively faster than men is that women had a comparative disadvantage at agriculture work, just as our theory predicts. To support their argument, Goldin and Sokoloff estimate that in 1820, women earned roughly 30% as men in the Middle Atlantic region, and roughly 37% as much as men in New England. By 1850, they estimate relative wages of 51% in the Middle Atlantic and 46% in New England. While numerous factors were at play in this period, the authors argue that their finding of rising female wages "is consistent with the observations of many contemporaries of the early nineteenth century who reported that the relative productivity (and wages) of women and children compared to adult men was low in the agriculture and traditional sectors of the pre-industrial northeastern economy (1982, page 759)."

As additional support for the comparative advantage theory, Goldin and Sokoloff provide evidence that women faced greater comparative disadvantage in farming in the North than in the South, and entered manufacturing to a much greater extent in the North than in the South. The difference in the comparative disadvantage of women stemmed from the types of farm work common in the two regions. In the North, where strength-intensive wheat farming was prevalent, women earned around one third as much as men in the 1820. In the South, where Cotton and Tobacco farming were most common, women earned around one half as much as men, as dexterity played a more important role in farming these crops. Just as the theory predicts, as the U.S. structural change progressed in the second half of the 19th century, Northern women entered into factory work to a much larger extent than those in the South.

6 Extension to Land as Fixed Factor of Agriculture Production

In this section we discuss why our mechanism applies more broadly, in particular if land is a fixed factor of agriculture production. The basic idea is that when economies have low overall productivity, subsistence needs force large fractions of workers to produce food on the fixed supply of farmland. This reduces agricultural productivity because of the decreasing returns to scale induced by the fixed factor. Thus, in the cross section of countries, differences in agricultural output per worker will be larger than the underlying productivity differences.

In future work we hope to explore two aspects of this idea further. First, we plan to extend our benchmark model to include land as a fixed factor of production in the agriculture sector. This will allow us to quantify the additional explanatory power that land adds to the baseline model with just worker specialization differences.

The second way we wish to explore this land idea further is to test the hypothesis that the quality of land used by the average agricultural worker in a developing country is lower than the quality of land used by agriculture workers in the developed world. This is a central prediction of the model where land as a fixed factor. Evidence from the FAO's Digital Soil Map of the World constructed by Wiebe (2003) suggests that indeed, farm workers in poor countries are on average using lower quality land. For example, just 5% of land actively farmed in Sub-Saharan Africa is deemed of high quality compared to over 30% for high-income countries. We plan to explore this data in more detail in future work.[8]

7 Extension to Trade

In this section we ask how allow the model's predictions would change once we allow for trade. We draw two conclusions. First, allowing for frictionless trade would introduce strongly counterfactual assumptions about shares of labor in agriculture and relative prices across countries. Thus, to be useful, any extension to allowing trade should have to include trade frictions and account for differences in relative prices across countries. Yet, if a model with trade is consistent with relative agriculture prices, we conjecture that such a model would have a modest effect on the *quantitative* nature of the model's predictions.

First consider a version of the model where each country has frictionless access to trade in world markets. Then the following is true.

Proposition 2 *Imagine that the rich and poor economies can trade frictionlessly on world markets at a relative food price p_a^W. Then the following must hold:*

$$\frac{Y_a^P/L_a^P}{Y_a^R/L_a^R} = \frac{Y_n^P/L_n^P}{Y_n^R/L_n^R} \quad \textit{and} \quad \frac{L_a^R}{L_a^R + L_n^R} = \frac{L_a^P}{L_a^P + L_n^P}. \tag{11}$$

Proposition 2 says that under frictionless trade, two things are true. First, the extent of specialization would be the same in both countries, and hence labor productivity differences between the rich and poor countries would be the same in agriculture and non-agriculture. Second, the shares of labor in agriculture would be equated across countries. Both are true because, under

[8] We thank Keith Wiebe for sharing his data with us.

a common relative price, the sector labor supply cutoff (5) is identical in both countries, and hence the composition of workers in each sector are identical as well.

But the prediction of labor shares being equal across countries is strongly counterfactual. As is well known, a substantially higher fraction of labor in poor countries is in agriculture than rich countries. In Section 2, for example, we cite evidence that the United States has just 2.8% of its labor in agriculture compared to over 78.3% of the labor in a country at the 10th percentile of the world per-capita income distribution. Thus, we conclude that treating economies as closed is the most sensible benchmark, as opposed to allowing countries access to frictionless trade on world markets.

Nevertheless, we conjecture that adding frictional trade would have a modest quantitative effect on the model's predictions in these three dimensions. The reason is that any reasonable model of trade would have to be in line with relative agriculture prices across countries, yet the baseline model is already consistent with the data in terms of prices. Thus we conjecture that the model's predictions for sector productivity differences would be changed little because – whether the model is an open economy or not – relative prices determine sectoral labor allocations, and sectoral labor allocations determine predictions about productivity.

8 Conclusion

We argue that cross-country productivity differences in agriculture are larger than in non-agriculture because of differences in the extent to which workers specialize in sectors in which they are relatively most able. In poor countries, virtually everyone works in agriculture, even though many of those workers have a comparative advantage that is *not* in farm work, but rather in non-agricultural tasks such as acting, teaching, or writing newspaper articles. In rich countries, in contrast, those remaining in agriculture are those who are relatively most productive at farm work. As a result, labor productivity differences are relatively larger in agriculture than the aggregate, and smaller in non-agriculture, even though countries differ only in general, sector-neutral, efficiency.

Our theory has new implications for the way economists think about agricultural productivity in the developing world. In contrast to other papers that emphasize barriers to efficient production in farming, we argue that low productivity in agriculture could represent the optimal response to low general efficiency in the face of subsistence food requirements. In this case it is optimal to employ many workers in agriculture who are less able in farm labor than other tasks. Concretely, our paper suggests that the source of low agriculture productivity might not be entirely found in the agriculture sector itself. It could, for example, be due to weak institutions, poor protection of property rights, or poor social infrastructure, as emphasized by a growing macroeconomics literature (e.g. Hall and Jones, 1999; Acemoglu, Robinson and Johnson, 2002).

A Model Appendix

A.1 Proof of Proposition 1

Let p_a^1, Y_a^1 and Y_n^1 be the equilibrium relative price and quantities in an economy with general efficiency A^1. Let $A^2 > A^1$, and denote by p_a^2, Y_a^2 and Y_n^2 the equilibrium of an economy with efficiency A^2.

Suppose that $p_a^2 = p_a^1$. Then by (5), each agent i chooses to work in the same sector in A^2 as in economy A^1. Thus output in each sector would be scaled up by a factor equal to the ratio of the efficiency terms: $Y_a^2/Y_a^1 = Y_n^2/Y_n^1 = A^2/A^1$. But by (6), we know that agents must demand a higher fraction of non-agriculture goods in economy A^2 than A^1. Thus $Y_n^2/Y_a^2 > Y_n^1/Y_a^1$. But this implies that $Y_n^2/Y_n^1 > Y_a^2/Y_a^1$, which is a contradiction. Thus $p_a^2 \neq p_a^1$.

The only way to be consistent with the agent solutions', (6), is for more agents to supply labor in the non-agriculture sector in economy A^2 than economy A^1. By (5), this occurs if and only if $p_a^2 < p_a^1$. ■

A.2 The Quantitative Implications of Alternative Parameter Values

In this section, we explore how the main parameters of interest influence our results and how these parameters of interest are disciplined by data. In particular, we show how the model's predictions change once we vary the dependence parameter ρ, the shape parameter controlling the agricultural ability distribution θ_a, and the shape parameter controlling the agricultural ability distribution θ_n. In all the calculations, we kept all other calibrated parameters at their baseline value and simply changed the parameter of interest. Panel A of Table 6 shows that

Table 6: Sensitivity Analysis

Panel A: The Effect of Dependence on the Model's Predictions

	Dependence Parameter ρ				
Prediction	0	1.03*	2	4	20
Percent Y/L explained in agriculture	88	78	06	-20	-40
Percent Y/L explained in non-agriculture	67	67	66	73	87
Ratio of average wage: $\bar{w}_a/\bar{w}_n$	0.79	0.77	0.66	0.54	0.39

Panel B: The Effect of Variation in Talent on the Model's Predictions

	Parameter θ_n				
Prediction	2	2.26*	4	6	8
Percent Y/L explained in agriculture	70	78	100	108	110
Percent Y/L explained in non-agriculture	72	77	48	39	36
Std deviation of log wages in Non-Agriculture	0.64	0.57	0.32	0.21	0.16
	Parameter θ_a				
Prediction	2	2.88*	4	6	8
Percent Y/L explained in agriculture	197	78	0.24	-6	-17
Percent Y/L explained in non-agriculture	41	67	81	92	97
Std deviation of log wages in Agriculture	0.61	0.46	0.36	0.26	0.20

increasing ρ decreases the predictive power of the model for agriculture productivity gaps yet leads to counterfactual implications for relative wage rate across sectors. When ρ ranges from its calibrated value of 1.03 to 2, we see that the model's explanatory power drops from 78% down to 6%.

Reducing ρ leads to counterfactual implications for wages. As ρ increases, the ratio of average wages in agriculture and non-agriculture decreases from 0.77 (as in the CPS data) down to 0.66. Thus, while our model's main predictions are sensitive to the value of ρ, average wage data by sector constrains the choice of ρ.

The results in Panel B show how the shape parameters θ_a and θ_n in the two sectors affect the model's explanatory power. As θ_a and θ_n increase, the model's explanatory power falls in sector the parameter is changed while the explanatory power in the opposite sector increases. For example, if θ_a increases from 2.26 to 6 then the model's explanatory power declines from to 78 percent to -6 percent. However, as one decreases θ_a, the models explanatory power increases for the opposite sector (non-agriculture in this case) from 67 percent to 92 percent when θ is set equal to 6. Also reported in Panel B are the implications of the model for the standard deviation of log wages in the sector with which the parameter was changed. Not surprisingly, when θ_a and θ_n are increased, the standard deviation of log wages in the respectively sector falls. This illustrates how variation in sector wage rates constrains our choice of θ in each sector.

B Data Appendix

- **GDP Per Worker** — This data is from the Penn World Table version 6.2. series "rgdpch".

- **Labor Share in Agriculture** — This data comes from Table A.3 in the FAO Statistical Yearbook 2004 online edition.

- **Agriculture Share in GDP** — This data comes from Table G.1 in the FAO Statistical Yearbook online edition.

- **Relative Agriculture Prices** — This data is derived from author's calculations with original data from the World Bank's 2005 International Comparison Program online database. The sector "agriculture" is defined to be food and non-alcoholic beverages, alcoholic beverages and tobacco, codes (1101 and 1102). "Non-agriculture" is defined as all individual consumption, code (11), gross fixed investment, code (15), minus food, non-alcoholic beverages, alcoholic beverages and tobacco.

- **U.S. Cross-sectional Wage Data** — We get our cross-sectional data on wages from the 2007 U.S. Current Population Survey (CPS). Following the study of U.S. wage inequality by Heathcote, Perri and Violante (2009), we take all individuals between 25 and 60 who have non-missing data on income and hours worked. Wages are before tax, and equal the sum of wage income plus two-thirds of business and farm income. We restrict the sample further to include only workers averaging at least 35 hours per week of work, and only those earning at least the Federal minimum wage. Wages are the total of wage, business and farm income before taxes. We include all individuals who are between 25 and 60, and any who didn't work at leas 1750 hours the previous year. We also drop any individual earning less than the federal minimum wage. Farmers are those whose occupational codes relate directly to agricultural production, fishing, forestry, or the raising of livestock.

- **U.S. Historical Relative Prices** — U.S. historical relative prices are from Historical Statistics of the United States Millennial Edition Online, Table Cc125-137 - Wholesale price indexes for historical comparisons, by commodity group: 1860 - 1990. Agriculture price is defined to be farm products and non-food is all commodities other than food products. As an alternative price series, we also explored using a series from Table Cc1-2, Consumer price indexes BLS based, in the denominator instead which yielded similar results. This alternative series is the analog to that used in Caselli Colmen (2005). To match up with observations on employment in Farming, observations corresponding with 1880, 1890, and 1900 are taken to be decade averages.

- **U.S. Historical Farm Population** — U.S. historical farm population are from Historical Statistics of the United States Millennial Edition Online, Table Da14-27 - Farmsnumber, population, land, and value of property: 1850 - 1997. This is taken to be a proxy for the share of employment in the United States in agriculture.

References

Acemoglu, Daron, Simon Johnson, and James Robinson. 2002. "Reversal of Fortune: Geography and Institutions in the Making of the Modern World Income Distribution." *Quarterly Journal of Economics* 117 (4): 1231–94.

Adamopoulos, Tasso. 2009. "Do Richer Countries Have Higher Distribution Margins?" Unpublished Manuscript, York University.

Allen, Robert. 1994. "Agriculture During the Industrial Revolution." Chapter 5 of *The Economic History of Britain Since 1700*, edited by Roderick Floud and Donald McCloskey, Volume 1, 96–122. Cambridge University Press.

Atkeson, Andrew, and Masao Ogaki. 1996. "Wealth-Vary Intertemporal Elasticities of Substitution: Evidence from Panel and Aggregate Data." *Journal of Monetary Economics* 38:507–534.

Caselli, Francesco. 2005. "Accounting for Cross-Country Income Differences." In *Handbook of Economic Growth*, edited by P. Aghion and S. Durlauf.

Caselli, Francesco, and Wilbur John Coleman II. 2001. "The U.S. Structural Transformation and Regional Convergence: A Reinterpretation." *Journal of Political Economy* 109 (3): 584–616.

Chanda, Areendam, and Carl-Johan Dalgaard. 2008. "Dual Economies and International Total Factor Productivity Differences: Channelling the Impact from Institutions, Trade and Geography." *Economica* 75:629–661.

Cherubini, Umberto, Elisa Luciano, and Walter Vecchiato. 2004. *Copula Methods in Finance*. John Wiley and Sons, Ltd.

Duarte, Margarida, and Diego Restuccia. 2009. "The Role of the Structural Transformation in Aggregate Productivity." *Quarterly Journal of Economics*. Forthcoming.

Eaton, Jonathan, and Samuel Kortum. 2002. "Technology, Geography and Trade." *Econometrica* 70 (5): 1741–1779.

Foster, Andrew D., and Mark R. Rosenzweig. 1996. "Comparative Advantage, Information and the Allocation of Workers to Tasks: Evidence from Agricultural Labor Markets." *Review of Economic Studies* 63 (3): 347–374.

Goldin, Claudia, and Kenneth Sokoloff. 1982. "Women, Children, and Industrialization in the Early Republic: Evidence from the Manufacturing Censuses." *Journal of Economic History* 42 (4): 741–774.

———. 1984. "The Relative Productivity Hypothesis of Industrialization: The American Case, 1820 to 1850." *Quarterly Journal of Economics* 99 (3): 461–487.

Gollin, Douglas, Stephen L. Parente, and Richard Rogerson. 2004. "Farm work, home work and international productivity differences." *Review of Economic Dynamics*, vol. 7.

Graham, Bryan S., and Jonathan R.W. Temple. 2006. "Rich Nations, Poor Nations: How Much Can Multiple Equilibria Explain?" *Journal of Economic Growth* 11, no. 1 (March).

Guvenen, Fatih, and Burhanettin Kuruscu. 2009. "A Quantitative Analysis of the Evolution of the U.S. Wage Distribution: 1970-2000." *NBER Macro Annual*. Forthcoming.

Hall, Robert E., and Charles I. Jones. 1999. "Why Do Some Countries Produce So Much More Output per Worker than Others?" *Quarterly Journal of Economics* 114 (1): 83–116 (February).

Heathcote, Jonathan, Perri Fabrizio, and Gianluca Violante. 2009, August. "Unequal we Stand: An Empirical Analysis of Economic Inequality in the United States, 1967-2006." Prepared for *Review of Economic Dynamics* Special Issue.

Heckman, James J., and Bo Honoré. 1990. "The Empirical Content of the Roy Model." *Econometrica* 58:1121–1149.

Heckman, James J., and G. Sedlacek. 1985. "Heterogeneity, aggregation, and market wage functions: an empirical model of self-selection in the labor market." *Journal of Political Economy* 93 (6): 1077–1125.

Kotz, Samuel, and Saralees Nadarajah. 2000. *Extreme Value Distributions*. Imperial College Press.

Postel-Vinay, Fabien, and Jean-Marc Robin. 2002. "Wage Dispersion with Worker and Employer Heterogeneity." *Econometrica* 70 (6): 2295–350.

Ramondo, Natalia, and Andres Rodríguez-Clare. 2009. "Trade, Multinational Production,and the Gains from Openness." Unpublished Manuscript, Penn State.

Restuccia, Diego, Dennis Tao Yang, and Xiaodong Zhu. 2008. "Agriculture and Aggregate Productivity: A Quantitative Cross-Country Analysis." *Journal of Monetary Economics* 55:234–250.

Rosenweig, Mark R., and Kenneth I. Wolpin. 1993. "Credit Market Constraints, Consumption Smoothing, and the Accumulation of Durable Production Assets in Low-Income Countries: Investments in Bullocks in India." *Journal of Political Economy* 101:223–244.

Roy, A. 1951. "Some Thoughts on the Distribution of Earnings." *Oxford Economic Papers* 3:135–46.

Townsend, Robert M. 1994. "Risk and Insurance in Village India." *Econometrica*, vol. 62 (3).

Vollrath, Dietrich. 2009. "How Important are Dual Economy Effects for Aggregate Productivity?" *Journal of Development Economics*, March, 325–334.

Wiebe, Keith, ed. 2003. *Land Quality, Agricultural Productivity, and Food Security*. Cheltenham, UK: Edward Elgar Publishing Limited.