Specialization, Agriculture, and Cross-Country Productivity Differences

David Lagakos

Arizona State University

Michael E. Waugh

New York University

August 2010

ABSTRACT

Cross-country 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 subsistence consumption needs prevent workers in poor countries from specializing in the sector in which they are most productive. We formalize our theory in a general-equilibrium Roy model in which the agents' preferences feature a subsistence consumption constraint. A parameterized version of the model predicts that output per worker gaps are substantially larger across countries in agriculture than non-agriculture even though countries differ only by a economy-wide efficiency term which affects the two sectors uniformly.

Email: lagakos@asu.edu, mwaugh@stern.nyu.edu. An earlier version of the paper circulated under the title "Specialization, Economic Development, and Aggregate Productivity." For helpful comments we thank Francesco Caselli, Berthold Herrendorf, Narayana Kocherlakota, Erzo Luttmer, Ellen McGrattan, Guillermo Ordoñez, Diego Restuccia, Andrés Rodríguez-Clare, Richard Rogerson, Jim Schmitz, Jon Vogel and seminar participants at ASU, Boston University, Clemson, the Federal Reserve Banks of Minneapolis and San Francisco, Notre Dame, Ohio State, Iowa, Syracuse, Toronto, UCLA, 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 Diego Restuccia.

1. Introduction

Cross-country labor productivity differences are much larger in agriculture and much smaller in non-agriculture than in the aggregate (Caselli (2005), Restuccia, Yang, and Zhu (2008)). Because developing countries have most of their workforce in agriculture, these sector productivity differences account for nearly all of aggregate productivity differences. This suggests that understanding productivity differences in agriculture is 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 economy-wide efficiency differences—which affect all sectors uniformly—combine with subsistence food consumption needs to prevent heterogeneous workers from specializing in the sector in which they are most productive.

The basic idea is that in poor countries, with low economy-wide efficiency, most of the workforce must be employed in the agriculture sector in order to satisfy subsistence needs. This is the "food problem" of Schultz (1953). Our insight is that if workers are heterogenous in their productivity in each sector, then many of those working in agriculture are those whose productivity advantage *is not* in agricultural work, but rather in non-agricultural tasks such as writing newspaper articles or teaching economics. In rich countries, with high economy-wide efficiency, those few workers remaining in agriculture are those whose productivity advantage *is* in agricultural work. As a result, measured productivity differences in agriculture are larger than the economy-wide efficiency differences. In non-agriculture the opposite occurs, and cross-country productivity differences in non-agriculture are smaller than in the aggregate.

Our theory has two main ingredients. First, preferences have a subsistence food requirement, and hence the income elasticity of demand for agricultural goods is less than one. Second, workers are heterogeneous in their productivity in the agriculture and non-agriculture sectors, and choose where to work. This is the Roy (1951) model of occupational choice embedded in a two sector general-equilibrium model. Countries differ only in a economy-wide efficiency term; preferences and the distribution of individual productivity are taken to be identical across countries to put focus on our mechanism.

Qualitatively, we show that the model can deliver productivity differences in agriculture that are larger than in the aggregate, and non-agriculture productivity differences that are smaller than in the aggregate. This result occurs when the individual productivity distribution is such that agents with a comparative advantage in a given sector also have an absolute productivity advantage in that sector.

To understand the quantitative importance of our theory we make flexible parametric assumptions on the distribution of individual productivity. These assumptions allow us to calibrate the the distribution parameters using simple moments from the distribution of wages in agriculture and non-agriculture. Preferences are standard and are disciplined by evidence on the importance of subsistence food needs in developing countries, and agriculture employment shares in developed countries.

Our main quantitative exercise varies economy-wide efficiency so that the model replicates a factor of 22 difference in aggregate productivity between the 90th and 10th percentile of the country income distribution, as observed in the data. We then compute the model's predictions for agriculture and non-agriculture productivity differences and compare them to data. In the data, the corresponding sector productivity gaps are a factor of 45 in agriculture and 4 in non-agriculture. Our model generates a factor of 35 difference in agriculture productivity and a factor 8 difference in non-agriculture. If the model had no predictive power, these sector productivity differences would equal the aggregate difference of 22. Thus, our model explains roughly half of the cross-country differences between agriculture and the aggregate, and more than half of the differences between non-agriculture and those in the aggregate. We show the model's quantitative effects depend in large part on the individual productivity dispersion across agents in the model, which are disciplined by the cross-sectional wage variation in U.S. data.

We find that our model is quantitatively consistent with other salient features of the data not targeted directly. In the cross-section of counties, our model generates much larger shares of employment in agriculture in poor countries than rich countries consistent with data. Our model predicts that the relative price of agricultural goods should be around three times higher in the poorest countries than the richest. We document that, in cross-country price data, the relative price of agriculture goods is at least twice as high in the poorest countries. In U.S. time series data, we compare our model's predictions to the same set of statistics as the U.S. developed over the past century and find our model's predictions are in line with the data. Finally, we show that the model's wage distribution closely resembles the observed wage distribution in the United States, and that plausible alternatives fail on this dimension.

To support our theory, we provide direct evidence of our mechanism using data on the composition of agriculture workers by sex and age across countries and over time. Under the assumption that women and children are at a comparative disadvantage in agriculture work (which is strength-intensive) compared to adult men (who tend to be stronger), our theory predicts that women and children should constitute a higher fraction of agricultural employment in poor countries than in richer countries. Using aggregate and census micro data from a large set of countries, we show that this is the case. Furthermore, our theory predicts that, during a period of structural transformation, women and children should leave agriculture at a faster rate than men. We cite evidence from Allen (1994) and Goldin and Sokoloff (1982, 1984) that this was the case during the development experiences of Britain and the United States.

Our work contribute to the large literature on understanding cross-country differences in per capita income and aggregate productivity (e.g. Klenow and Rodríguez-Clare (1997), Hall and Jones (1999), Parente and Prescott (2000), and Caselli (2005)). The accounting results of Caselli (2005) and Restuccia, Yang, and Zhu (2008) suggest that low aggregate productivity is driven by large fractions of workers working in the relatively unproductive agriculture sector. Our work, in contrast, suggests that lower measured productivity in agriculture is a *consequence* of aggregate factors, such as weak institutions or poor social infrastructure (Hall and Jones (1999); Acemoglu, Johnson, and Robinson (2001)). Our model can interpreted as a theory of how economy-wide efficiency differences across countries–affecting all sectors of the economy in the same way–are amplified into even larger agriculture productivity differences.

The policy implications of the two views are quite different. According to the view that low aggregate productivity is driven by low agriculture productivity, fixing agriculture is the most important step in development. In contrast, an implication of our theory is that general improvements in technology, social infrastructure, or institutions are the key to improving living standards, and that agriculture productivity will follow.

2. Motivating Evidence

In this section, we show that cross-country labor productivity differences are large in agriculture and smaller in non-agriculture, relative to aggregate differences. We then discuss how our paper fits into and contributes to the literature on the role of agriculture in understanding aggregate productivity.

Table 1 reproduces data from Caselli (2005), who constructs Purchasing Power Parity (PPP) adjusted measures of output per worker in agriculture and non-agriculture sectors for 79 countries in 1985. His calculations combine PPP GDP per worker data from the Penn World Tables with PPP agriculture value added per worker data from the Food and Agriculture Organization (FAO)(Rao (1993)).[1] Panel A reports that the ratio of agricultural output per

[1] In independent work, Restuccia, Yang, and Zhu (2008) arrive at findings similar to those in Table 1.

worker in the 90th and 10th percentiles of the income distribution is 45, compared to just four in non-agriculture. As a frame of reference, the ratio for aggregate output per worker is 22. Thus, across countries, labor productivity in agriculture is roughly twice as variable relative to labor productivity in the aggregate, and non-agriculture is roughly one-fifth as variable as the aggregate.[2]

Table 1: Labor-Productivity and Labor Allocations

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)

Simple accounting exercises show that the divide between agriculture and non-agriculture accounts for aggregate productivity differences. Caselli (2005) computes the hypothetical ratio of aggregate output per worker between the 90th and 10th percentile countries by equalizing agricultural productivity to the U.S. level in all countries. He finds that it is a factor of 1.6, down from the actual factor of 22! In other words, cross-country income inequality would be eliminated if agriculture productivity were equalized to the U.S. level. In a related exercise, Caselli (2005) computes the hypothetical ratio of aggregate output per worker between the 90th and 10th percentile countries, assuming that agriculture's share of labor in all countries is equal to 2.8 percent of workers in agriculture, as in the 90th percentile country (Panel B of Table 1). He finds that the hypothetical ratio of aggregate output per worker is a factor 4.2, which is again far lower than the actual 22 in the data. These accounting exercises suggest that the divide between agriculture and non-agriculture is intimately related

[2]These findings relate to a long history of research in macroeconomic development, dating back to Schultz (1953), which has found that value added per worker is much higher in non-agriculture than in agriculture. Recently, Gollin, Parente, and Rogerson (2004) examine data from a sample of countries in 2000 and find that non-agriculture value added per worker is far higher than agricultural value added per worker, often by a factor of 10.

to understanding aggregate productivity differences.

Research on *why* productivity differences are larger in agriculture than in the aggregate has advanced mainly in one direction. Specifically, researchers, have asked whether the lack of productive inputs in agriculture or barriers to efficient agriculture production are responsible for agriculture's low productivity. Chanda and Dalgaard (2008) and Vollrath (2009) construct physical and human-capital measures by sector and find that neither can explain agriculture productivity differentials. They conclude that low agriculture labor productivity reflects low agriculture total factor productivity (TFP). Restuccia, Yang, and Zhu (2008) argue that barriers to the use of intermediate inputs in agriculture and barriers in the labor market lower agriculture productivity in poor countries.

Our paper proposes a different view. Rather than sector-specific distortions leading to sector productivity differences, we argue that these differences arise *because* of economy-wide efficiency differences (e.g., institutions, social infrastructure, and economy-wide barriers to technology adoption), which lead to cross-country differences in average productivity of labor across sectors.[3]

3. Model of Agricultural and Non-Agricultural Productivity

In this section, we formalize our model economy and then characterize its equilibrium. The model predicts, under circumstances that we describe, that lower economy-wide efficiency results in even larger labor productivity differences in agriculture and smaller labor productivity differences in non-agriculture. Two ingredients in our model deliver this outcome. First, households have preferences with a subsistence food requirement. Second, workers are heterogeneous in their productivity in the two sectors.

3.1. Households

There are measure one of agents, indexed by i, who differ in productivity, as explained below. Preferences are given by

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

[3]Papers such as Gollin, Parente, and Rogerson (2004) and Herrendorf and Teixeira (2009) combine both views by arguing that barriers to capital accumulation (in either sector), in addition to economy wide and sector specific distortions, lead to lower measured agriculture productivity. Graham and Temple (2006) proposes an altogether different mechanism by arguing that increasing returns in non-agriculture and decreasing returns in agriculture lead to two equilibria, with the poor countries in the bad one with most workers in the least productive sector.

where c_a^i is agricultural good (food) consumption, c_n^i is non-agricultural good consumption, $\bar{a}$ is a parameter representing a subsistence agriculture requirement, and ν governs the relative taste for non-agriculture consumption. These "Stone-Geary" preferences ensure that food is a necessity, and consumed in higher amounts relative to the non-agricultural good when overall expenditure is lower.

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 "individual productivities" $\{z_a^i, z_n^i\}$, which represent the efficiency of one unit of labor in sectors a and n. Individual productivities are drawn from a distribution $G(z_a, z_n)$ with support on the positive reals. Agents earn labor income y^i, which is described in more detail below. The budget constraint is

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

where p_a is the relative price of agriculture, 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 with its own production function. Both sector technologies are freely available and operated by competitive entrepreneurs. The technologies are given by

$$Y_a = AL_a \quad \text{and} \quad Y_n = AL_n \tag{3}$$

in agriculture and non-agriculture, where A is exogenous and captures "economy-wide efficiency" of production, and L_a and 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 choosing to work in agriculture and non-agriculture. The sector aggregate labor inputs L_a and L_n are defined as

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

and represent the sum of all individual productivity employed in the sectors. The total number of workers' in each sector is defined as

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

3.3. Optimization and Equilibrium

An equilibrium of the economy consists of a relative agriculture price, p_a, wages per efficiency unit of labor in each sector, w_a and w_n, and allocations for all agents, such that all agents optimize and labor and output markets clear. Measured labor productivity in equilibrium is denoted by Y_a/N_a in agriculture and Y_n/N_n in non-agriculture, and represent the physical quantity of output produced per worker in each sector.

Agents take prices and sector wages as given when they optimize. The problem for an agent is first to pick which sector to work in, and then to maximize her utility (1) subject to its budget constraint (2). Because of competition in production markets, the wages per efficiency unit of labor are

$$w_a = p_a A \quad \text{and} \quad w_n = A$$

in the agricultural and non-agricultural sectors. A simple cutoff rule in *relative* individual productivity, or comparative advantage, determines the optimal occupational choice for each agent. Working in non-agriculture is optimal for agent i if and only if

$$\frac{z_n^i}{z_a^i} \geq p_a. \tag{4}$$

Thus, the agents that enter non-agriculture are those whose productivity is sufficiently high relative to their productivity in agriculture. Labor income under the optimal sector choice is defined as $y^i \equiv \max\{z_a^i w_a, z_n^i w_n\}$.

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

$$c_a^i = \frac{y^i + \bar{a} p_a \nu}{p_a(1+\nu)} \quad \text{and} \quad c_n^i = \frac{\nu(y^i - \bar{a} p_a)}{1+\nu}. \tag{5}$$

Due to the subsistence agriculture requirement, agents consume relatively more agricultural goods when their labor income is lower. The lower is ν, the higher the ratio of agriculture to non-agriculture consumption.

4. Qualitative Features

In this section we characterize qualitative features of the model, and we provide an analytical example (and special case of our quantitative model) to develop intuition about how the model works.

4.1. Relative Price of Agriculture is Higher in Poorer Economies

In equilibrium, the relative price of agriculture is decreasing in A. As we demonstrate below, this is an important component of how economy-wide efficiency affects labor productivity in the two sectors.

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 lower in the rich economy: $p_a^R < p_a^P$.*

The intuition is that a higher price of agricultural goods is needed in the poor economy in order to induce workers to work in the agriculture sector. To see this, let p_a^R be the equilibrium relative price in rich economy. If this were the equilibrium price in the poor economy as well, then, by (4), the sector labor-supply cutoffs would be the same in both countries, hence, so would the share of workers in agriculture. But because of the subsistence agriculture requirement, the poorer economy demands a much larger fraction of agricultural goods. Thus there would be excess demand for food in the poor economy. It follows that the relative price of agriculture could not be the same in the two economies, and in fact must be higher in the poor economy than the rich economy.

4.2. Individual Productivity Distribution and Sectoral Productivity Differences

Proposition 2 describes conditions on the individual productivity distribution that are sufficient for economy-wide efficiency differences to lead to larger differences in agriculture labor productivity and smaller differences in non-agriculture labor productivity.

Proposition 2 *Consider two economies with efficiency terms A^R and A^P such that $A^R > A^P$. Let the individual productivity distribution be such that $E(z_a|z_a/z_n > x)$ and $E(z_n|z_n/z_a > x)$ are increasing in x. Then, equilibrium sector labor productivities are such that*

$$\frac{Y_a^R/N_a^R}{Y_a^P/N_a^P} > \frac{A^R}{A^P} \quad \textit{and} \quad \frac{Y_n^R/N_n^R}{Y_n^P/N_n^P} < \frac{A^R}{A^P}.$$

The first part of the proposition says that agriculture productivity differences are larger than A differences if expected agriculture individual productivity is higher for agents with a greater comparative advantage in agriculture. The reason is as follows. As A rises, the relative price of agriculture falls, and only agents with a greater comparative advantage in agriculture (i.e., higher z_a/z_n ratio) choose to work in agriculture. Agriculture productivity, overall, increases if expected agriculture individual productivity is higher for agents with a stronger comparative advantage. The second part says that non-agricultural productivity differences are smaller than A differences if agents with a greater comparative advantage in non-agriculture have a higher expected individual productivity in that sector.

At least one of Proposition 2's conditions on the individual productivity distribution must hold (see Heckman and Honoré (1990)). Thus, at the very least, our theory qualitatively delivers productivity differences in one sector that differ from the aggregate differences in a way consistent with the data. Of course, it can also explain the patterns of both sector productivity differences. We now turn to an example where both conditions on the individual productivity distributions are satisfied, and in which simple analytical expressions help provide intuition for how the model works.[4]

4.3. Analytical Example: Independent Fréchet Individual Productivities

In this section, we analytically illustrate the mechanics of the model assuming independent Type II extreme value distributions, or Fréchet distributions over individual productivity. This is also a special case of the individual productivity distribution used for quantitative analysis in Section 5.

Assumption 1 *Let z_a and z_n be drawn independently from Fréchet distributions:*

$$G(z_a) = e^{-z_a^{-\theta}} \quad \textit{and} \quad G(z_n) = e^{-z_n^{-\theta}}.$$

The parameter θ controls the dispersion of individual productivity in each sector, with a smaller θ implying more individual productivity dispersion across individuals and a higher θ meaning less dispersion.[5]

This distributional assumption conveniently relates employment shares in agriculture, the relative price of agriculture, and the individual productivity-dispersion parameter θ. To see this, the equilibrium share of workers in agriculture is

$$\pi_a = \text{Prob}\left\{Az_n^i \leq p_a A z_a^i\right\} = \frac{1}{p_a^{-\theta} + 1}. \tag{6}$$

[4]Other restrictions on the individual productivity distributions for which both conditions of Proposition 2 hold are those with log individual productivities that are independent across sectors and distributed log-concave in each sector. Prominent examples are Normal, Pareto and Uniform distributions. However, none has the analytic tractability productivity of independent Fréchet distributions that we focus on below.

[5]This distribution has been used by Eaton and Kortum (2002) and others to analytically solve multi-country Ricardian models of international trade. To our knowledge, we are the first to exploit the analytical properties of this distribution to study the Roy model.

Proposition 1's effects are seen in equation (6): As the relative price of agriculture increases the share of workers in agriculture also. Furthermore, the responsiveness of the share of workers in agriculture to the relative price depends on the individual productivity-dispersion parameter θ. Manipulating equation (6) yields a log-linear relationship in the ratio of the agriculture to non-agriculture worker shares (π_n) and the relative price of agricultural goods:

$$\log(\pi_a/\pi_n) = \theta \log(p_a). \tag{7}$$

The elasticity of relative employment shares between agriculture and non-agriculture to the relative price of agriculture is exactly θ. Intuitively, with a low θ, meaning high productivity dispersion across workers, large changes in the relative price of agriculture are needed to induce workers to switch sectors. On the other hand, a higher θ, meaning small individual productivity dispersion, implies that only small changes in the relative price of agriculture are needed to induce workers to switch sectors. Equation (7) is consistent with this intuition as to how individual productivity dispersion matters.

Both conditions of Proposition 2 hold in this example. That is, expected worker productivity in a sector is larger when its workers have a greater comparative advantage in that sector. In particular, expected individual productivity of workers in each sector is

$$E(z_a|z_a/z_n > 1/p_a) = \gamma\{(1/p_a)^\theta + 1\}^{\frac{1}{\theta}}, \text{ and } E(z_n|z_n/z_a > p_a) = \gamma\{p_a^\theta + 1\}^{\frac{1}{\theta}}, \tag{8}$$

where the constant γ is the Gamma function evaluated at $(\theta-1)/\theta$.[6] In the agriculture sector, the inverse of the relative agriculture price controls the degree of comparative advantage. In the non-agriculture sector, it's the relative agriculture price. Equation (8) shows that an increase in either leads to an increase in expected individual productivity in the respective sector. Thus, by Proposition 2, differences in A lead to larger productivity differences in agriculture and smaller ones in non-agriculture. We formalize this as the following Proposition.

Proposition 3 *Consider two economies with efficiency terms A^R and A^P such that $A^R > A^P$, and let Assumption 1 hold. Then, the ratios of sector labor productivities are*

$$\frac{Y_a^R/N_a^R}{Y_a^P/N_a^P} = \frac{A^R}{A^P}\left(\frac{(1/p_a^R)^\theta + 1}{(1/p_a^P)^\theta + 1}\right)^{\frac{1}{\theta}} > \frac{A^R}{A^P}$$

and

$$\frac{Y_n^R/N_n^R}{Y_n^P/N_n^P} = \frac{A^R}{A^P}\left(\frac{(p_a^R)^\theta + 1}{(p_a^P)^\theta + 1}\right)^{\frac{1}{\theta}} < \frac{A^R}{A^P}.$$

[6] The Gamma function is defined as $\Gamma(\xi) = \int_0^\infty y^{\xi-1}e^{-y}dy$ where $\xi > 0$.

The terms in parentheses are the ratios of average sector individual productivities conditional on working in the sector. Since the relative price of agriculture is higher in the poor economy, it follows that average agriculture individual productivity is relatively higher in the rich economy, and average non-agriculture individual productivity is relatively lower in the rich economy. Thus, labor productivity in agriculture is more variable, and productivity in non-agriculture is less variable than the underlying A differences.

Dispersion in individual productivity controls the magnitude of the sector productivity difference from the aggregate. Combining equations (6) and (8), relative productivity differences in agriculture across a rich and poor economy are

$$\frac{Y_a^R/N_a^R}{Y_a^P/N_a^P} = \left(\frac{\pi_a^P}{\pi_a^R}\right)^{\frac{1}{\theta}} \left(\frac{A^R}{A^P}\right) \geq \frac{A^R}{A^P}. \tag{9}$$

Equation (9) shows how a lower θ, meaning higher in individual productivity dispersion, leads agriculture productivity to be larger than the aggregate since $\pi_a^R < \pi_a^P$ in equilibrium. As θ approaches infinity, the ratio of agriculture productivity converges downward toward the aggregate productivity ratio.

A similar argument illustrates that non-agriculture productivity differences are smaller than the difference in A, with the magnitude of the difference depending upon individual productivity dispersion. Relative productivity differences in non-agriculture across a rich and poor economy are

$$\frac{Y_n^R/N_n^R}{Y_n^P/N_n^P} = \left(\frac{\pi_n^P}{\pi_n^R}\right)^{\frac{1}{\theta}} \left(\frac{A^R}{A^P}\right) \leq \frac{A^R}{A^P}. \tag{10}$$

Because $\pi_n^R > \pi_n^P$, a lower θ leads to non-agriculture productivity smaller than the aggregate, with the differences converging towards the aggregate as θ increases.

5. Quantitative Analysis

While the example above is useful in illustrating how the model works, several restrictive features limit its use in quantitative analysis—in particular, its assumptions of independence of sector-productivity draws across individuals, and productivity dispersion across individuals that is the same in each sector. We now turn to a richer model of the individual productivity distribution that relaxes both of these restrictions. We then use the model to ask whether exogenous economy-wide efficiency differences can lead to much larger productivity variation in agriculture than in non-agriculture.

5.1. Dependent Fréchet Individual Productivity Distribution

We set the joint distribution of individual productivities as

$$G(z_a, z_n) = C[F(z_a), H(z_n)],$$

$$\text{where} \qquad F(z_a) = e^{-z_a^{-\theta_a}} \text{ and } H(z_n) = e^{-z_n^{-\theta_n}},$$

$$\text{and} \qquad C[u, v] = \frac{-1}{\rho} \log\left\{1 + \frac{(e^{-\rho u} - 1)(e^{-\rho v} - 1)}{e^{-\rho} - 1}\right\}.$$

The function $C[F(z_a), H(z_n)]$ is a *Frank copula*, which allows for dependence between draws from distributions $F(z_a)$ and $H(z_n)$.[7] The parameter $\rho \in (-\infty, \infty)\backslash\{0\}$ determines the extent of dependence, with a positive (negative) value of ρ representing positive (negative) dependence between the draws.[8] The marginal distributions themselves are Fréchet, with dispersion parameters θ_a and θ_n. The lower are θ_a and θ_n, the higher is the variation in individual productivity in agriculture and non-agriculture.

This parameterization introduces two additional dimensions of richness relative to the example in Section 4.3. First, individual productivity draws are no longer independent across sectors. This allows for characteristics—such as intelligence or entrepreneurial ability—that make a worker more able in both types of activity. Second, dispersion in individual productivity is no longer the same in each sector. Since non-agriculture work is a stand-in for many different types of economic activities, one might expect that individual productivity dispersion is larger in non-agriculture than in agriculture. This parameterization allows for this possibility.

We chose this functional form for our quantitative analysis for three main reasons.[9] First, it allows for a transparent calibration of the distribution parameters, while also allowing for dependence and differing sector individual-productivity dispersion. As we show in the following section, the three parameters of the distribution (θ_a, θ_n and ρ) are disciplined by

[7] See Nelsen (2006). A copula is a function that allows for the creation of multivariate distributions out of arbitrary univariate distributions. This choice of the Frank copula generates dependence between draws that is not systematically stronger when closer to the right or left tails of the distribution—i.e. it's radially symmetric. Other copula, such as the Clayton or the Gumbel copula, do not have this feature.

[8] In the case of independence, the distribution is defined as $G(z_a, z_n) = e^{-z_a^{-\theta_a}} e^{-z_n^{-\theta_n}}$.

[9] Individual productivity distributions in the Roy model cannot be identified from cross-sectional wage data without making assumptions about the functional form of the distributions (see Heckman and Honoré (1990)). Because one observes only the maximum of each agents' draws, but not both draws themselves, if individual productivity distributions are allowed to take on an arbitrary form, there are many distributions that can generate a given set of observations on wages and sector choices by individuals.

three moments in a single cross-section of wages: the standard deviation of log wages in a sector and relative average wages across sectors.[10]

Second, the choice of Fréchet distributions for individual productivity 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.[11] Thus, the draw z_n^i, for example, can be thought of as the maximum of household i's individual productivity draws in a large set of distinct non-agricultural tasks. A similar interpretation can be given to z_a^i.

Third, Fréchet distributions and the Frank copula yield sector and overall wage distributions that resemble their empirical counterparts, as we demonstrate below. In particular, the model delivers wage-distribution tails that mimic the data closely, a dimension along which other distributions fail. For example, we find that a version of our model with log-normal sector individual-productivity distributions generates tails that are too thin compared to the data.[12]

5.2. Calibration of Individual Productivities

To calibrate the individual-productivity distribution parameters, our strategy uses cross-sectional wage data from the United States. Formally, we jointly calibrate 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. While all three parameters are jointly determined, each has an intuitive relationship with one of the moments picked, which we discuss below.

Intuitively, the θ_a and θ_n terms determine the variation in individual productivity across individuals, with higher θ's resulting in lower variation in individual productivities. Because wages are set equal to the value of marginal products, variation in individual productivity maps into variation in wages across agents. Thus, variation in agriculture and non-agriculture wages are key in disciplining the parameters θ_a and θ_n.

The dependence parameter ρ is informed 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 Figures 1(a) and 1(b). Both figures plot a simulation with each point in the plane

[10] In Section 9, we discuss how our results would change assuming differing individual-productivity distributions across countries.

[11] 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.

[12] See Appendix C. Heckman and Sedlacek (1985) also conclude that log-normal individual-productivity distributions lead to counterfactual predictions for the tail of the wage distribution.

depicting the two wages that a household can receive. Figure 1(a) has independent draws and Figure 1(b) has positive dependence in draws. If a household is to the left of the 45^o line, then it is optimal to work in agriculture, and vice versa. In both figures, only two percent of the workers chose agriculture, as in U.S data, and the θ's are calibrated to hit the standard deviation of log wages in agriculture and non-agriculture in our data. Also plotted is the log mean wage of workers in agriculture and non-agriculture across the two scenarios.

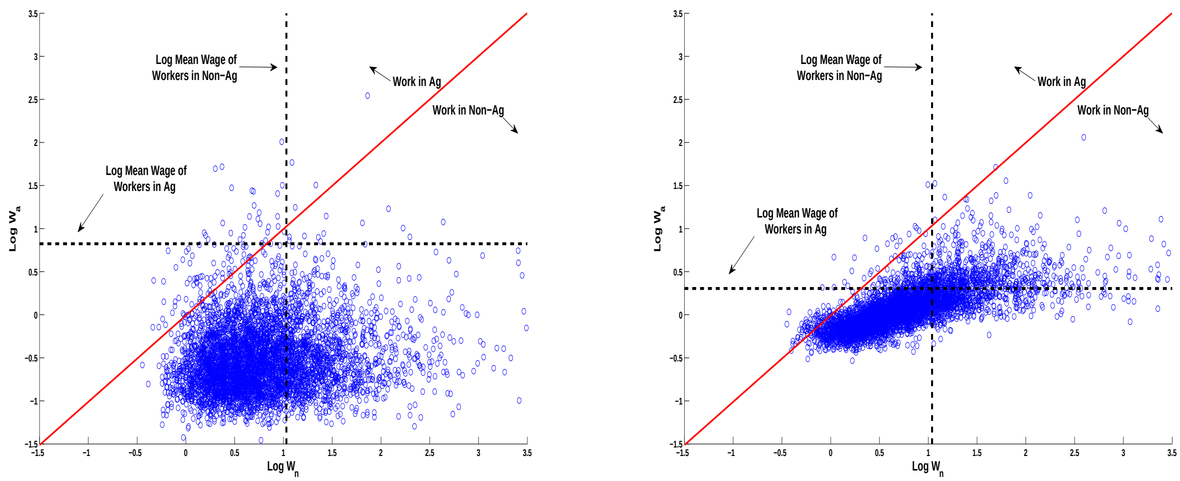

(a) Independent individual productivity.

(b) Positive Dependence in individual productivity.

Figure 1: Wage Offers under Independence and Dependence of Individual Productivity

The difference to notice between Figures 1(a) and 1(b) is that workers in agriculture earn, on average, lower wages (compared to non-agriculture) in the model with dependence relative to the model with independence. Mechanically, this can be seen in the way the horizontal dashed line (indicating the log mean wage in agriculture) shifts down across the two, while the vertical dashed line (indicating the log mean wage in non-agriculture) stays essentially the same across the two scenarios. Intuitively, higher dependence means that most of the low-draw workers end up in the sector with lower individual productivity variance—in this case, agriculture (as we discuss below). Because of this, a higher value of ρ leads to a lower ratio of average wages in agriculture to non-agriculture.

Our cross-sectional wage data comes from the U.S. Current Population Survey (CPS) for 2007. Following the study of U.S. wage inequality by Heathcote, Perri, and Violante (2009), our sample includes all individuals between ages 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 variation, which will lead to more conservative variation in individual productivity and, hence, more conservative results. We calculate that the standard deviation of log wages for agriculture and non-agriculture workers are 0.46 and 0.57. The average wage in agriculture is 0.68 times the average wage in non-agriculture.[13]

These moments imply parameter values of $\theta_a = 4.08$, $\theta_n = 2.27$ and $\rho = 2.90$. The estimates of θ_a and θ_n mean that there is more variation in individual productivity in non-agriculture work than in agricultural work, which seems reasonable given that non-agriculture work encompasses more types of economic activities. Since ρ itself is hard to interpret, we computed the implied linear correlation coefficient to be 0.39. This suggests that there is a substantial amount of positive correlation in sector individual productivities—i.e., if an agent is productive in one sector, she is likely to be productive in the other.

5.3. Calibration of Preference Parameters

For the preference parameters, we pick ν to match the fraction of workers in agriculture from U.S. data, which is just below two percent. The resulting parameter implies a 0.1 percent long-run expenditure share on food consumption. This is in line with other values used in the literature; Restuccia, Yang, and Zhu's (2008) model implies a long-run expenditure share of 0.46 percent; Gollin, Parente, and Rogerson (2004) implies a value of 0.3 percent; and in Gollin, Parente, and Rogerson (2007), the long-run expenditure share is zero.

We pick $\bar{a}$ to match a subsistence consumption need of 34 percent of average income in a model country with 7.5 percent of the U.S.'s per capita GDP. This is consistent with the independent estimates of subsistence agriculture consumption requirements in Rosenweig and Wolpin (1993), and Atkeson and Ogaki (1996), both of which use panel data from a sample of rural households in India (which had 7.5 percent of the U.S. per capita GDP in 1984, the year their data were collected).

5.4. Quantitative Predictions for Sector Productivity Differences

To explore the quantitative implications of our model, we perform the following experiment. Beginning with a value of economy-wide efficiency A normalized to one for the benchmark economy (calibrated to the U.S.), we lower A to match GDP per worker for a country in the 90th percentile of the income distribution, and then for a country in the 10th percentile. In

[13] If the cost of living is lower in rural areas, then the ratio of real average wages would be higher. This implies that we would infer a lower ρ, and hence, our results would be strengthened. See Section 9 for the sensitivity of our results to this parameter.

each case, we compute labor productivity in agriculture and non-agriculture, and compare the difference in sectoral output per worker across the two countries.

Table 2: Labor Productivity Differences

Ratio of 90th-10th Percentile			
Sector	Data	Model	Percent Explained
Agriculture	45	34.5	54
Aggregate	22	22	-
Non-Agriculture	4	8.2	77

Table 2 shows the model's predictions for the ratio of the 90th to the 10th percentile of countries in the model and data. The differences in aggregate output per worker is a factor of 22 in the model and data by construction.[14] If all workers were identical, the model would predict that agricultural and non-agricultural productivity ratios would be equal to the aggregate ratios. Instead, the model predicts that agriculture output per worker differences should be a factor of 34.5, and that non-agriculture differences should be a factor of 8.2. In the data, these ratios are a factor of 45 and four. The third column of the table shows that this corresponds to the model explaining 54 percent of the difference between the agricultural productivity ratio and the aggregate ratio, and 77 percent of difference between the non-agricultural ratio and the aggregate ratio.[15]

While our model does well in explaining differences between the 90th- and 10th- percentile countries, it does less to explain the differences between the 90th-percentile countries and those at intermediate income levels. Table 3 illustrates the model's prediction for the 90th-50th ratio and 90th-25th ratios. In the latter case, aggregate productivity in the model and data differs by a factor of 9.4, again by construction. The model predicts a factor of 16.1 in agriculture and 7.2 in non-agriculture, compared to 31.1 and 2.7 in the data. Thus, the model explains 31 and 33 percent of the agricultural and non-agricultural productivity differences relative to the aggregate. While still large, this is substantially lower than the 90th-10th percentile ratio.

In the 90th-50th case, the model explains even less. 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 4.3 and three, compared to 11.1 and 1.9 in the data. This amounts to explaining just 15 and eight percent of the sector differences compared to aggregate differences.

[14] Aggregate output per worker is expressed as GDP per worker at Gheary-Khamis international prices.

[15] Percent explained is measured as the predicted sector productivity ratio minus the aggregate ratio divided by the actual sector productivity ratio minus the aggregate ratio. For agriculture: (34.5-22)/(45-22).

Table 3: Labor Productivity Differences – Intermediate Income Levels

Sector	Data	Model	Percent Explained
Ratio of 90th.-25th Percentile			
Agriculture	31.1	16.1	31
Aggregate	9.4	9.4	-
Non-Agriculture	2.7	7.2	33
Ratio of 90th-50th Percentile			
Sector	Data	Model	Percent Explained
Agriculture	11.1	4.3	15
Aggregate	3.1	3.1	-
Non-Agriculture	1.9	3.0	8

Why does the model explain less of the sector productivity difference between rich and middle income countries? The shares of labor in agriculture across these countries are not as different; thus, average worker productivity differences have less room to be important. Consider the case of the 90th-50th differential. The share of workers in agriculture in the 50th-percentile country is nine percent, compared to three percent in the 90th-percentile country. Thus, the average productivity of workers in agriculture is only slightly lower than in the 50th-percentile country. In contrast, in the 10th-percentile country, 78 percent are in agriculture, so the average worker has substantially lower productivity than the average agricultural worker in the 90th-percentile country. Analytically, our example in Section 4.3 illustrates this point in Equations (9) and (10): The model's explanatory power is larger when labor shares differ greatly, as they do between rich and poor countries, but not between rich and middle-income countries.

5.5. Other Cross-Country Implications

In this section, we argue that the model's other quantitative predictions are consistent with the relevant macroeconomic facts from the cross-section of countries.[16]

We find that the shares of labor in agriculture in the model are consistent with cross-country data. Figure 2 plots data on the percent of employment in agriculture against PPP GDP per worker data for the year 2000. The country in the 10th percentile in GDP per worker has an employment share in agriculture of 78 percent. In contrast, the country in the 90th

[16] See Appendix B for more detail on our data sources.

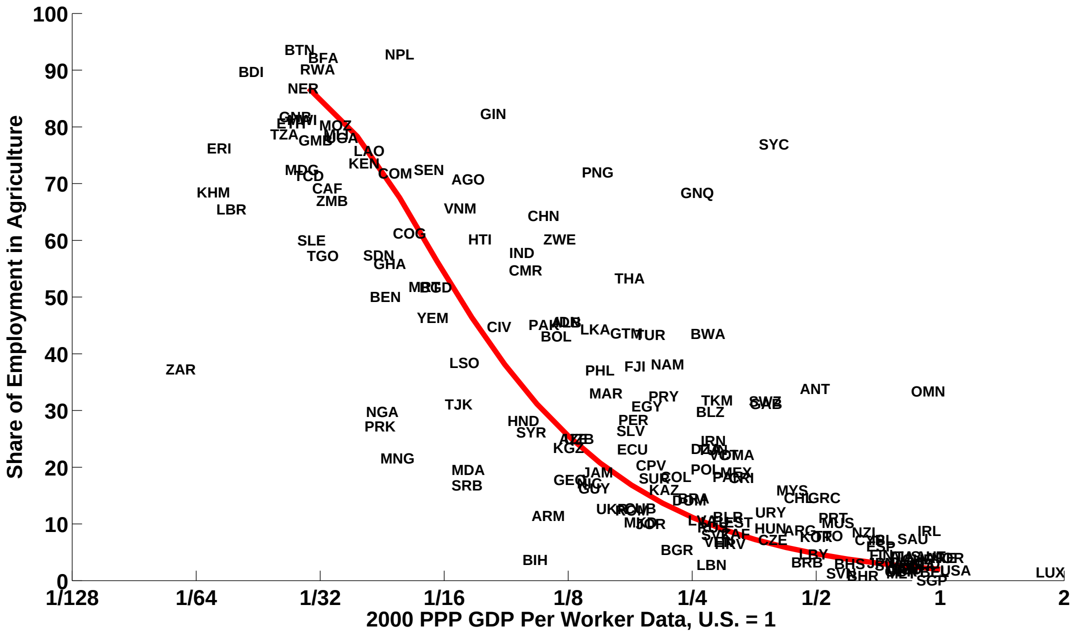

Figure 2: Share of Employment in Agriculture, Data and Model

percentile of the income distribution has a share of three percent. The model's predictions are generated by varying A across countries and are plotted as the solid line in Figure 2. The model predicts that a country in the bottom 10th percentile in GDP per worker should have around 80 percent of workers in agriculture, and that the percent declines with increases in GDP per worker in a way consistent with the data.

Figure 3 plots data on the relative price of agriculture versus PPP GDP per worker data as well as our model's prediction. Our data on relative agriculture prices are constructed using 2005 data from the International Comparison Programme (ICP); Appendix B provides the complete details. Figure 3 shows that relative agriculture prices systematically decline with increases in GDP per worker.[17] The ratio of relative prices between countries in the 90th and 10th percentiles of GDP per worker is 2.3. The solid line in Figure 3 plots our model's prediction. Relative agriculture prices systematically decline with increases in GDP per worker, with the ratio between the 90th and 10th percentiles being 3.3.

One concern is that the data are based on the prices that consumers pay for goods, not the

[17]This fact is consistent with previous studies of variation in cross-country relative prices—e.g., Summers and Heston (1991), Jones (1994), Restuccia and Urrutia (2001), and Hsieh and Klenow (2007). In particular, Herrendorf and Valentinyi (2009), show that when partitioning ICP goods into agricultural and non-agricultural goods, the relative price of agriculture is higher in poor countries. They also show that partitioning goods into tradeable and non-tradeable goods implies a higher relative prices of tradeables in poor countries, and partitioning goods into consumption and investment goods implies a higher relative price of investment goods in poor countries.

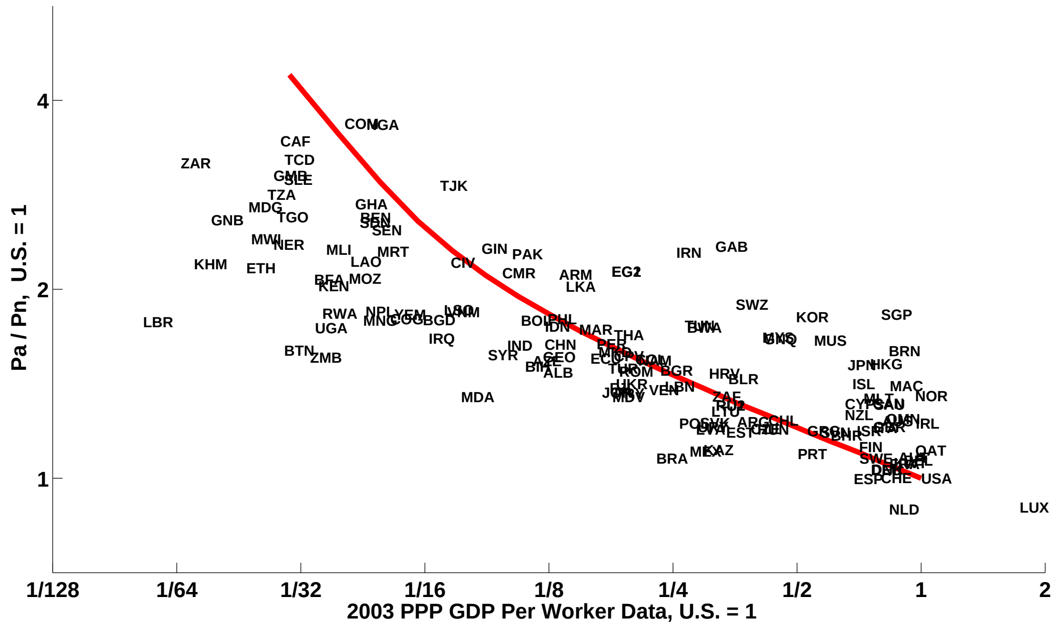

Figure 3: Relative Price of Agricultural Goods, Data and Model

price that producers receive. This distinction would reflect distribution margins that are not in the model. If distribution margins vary systematically with level of development, then the relationship in Figure 3 may not reflect differences in relative agriculture-producer prices. To address this concern, we examined relative agriculture-price data using producer prices – from a smaller sample of countries – constructed by Restuccia, Yang, and Zhu (2008). They used 1985 FAO agriculture-price data from Rao (1993) for the same 79 countries for which we have agriculture value added per worker data outlined in Table 1. We find that, by these measures, relative agriculture prices systematically decline with increases in GDP per worker—as our model predicts—and, in fact, the relationship is even stronger than for consumer prices.[18]

5.6. Implications for U.S. Cross-Sectional Distribution of Wages

In our calibration, we made parametric assumptions regarding the distribution of individual productivities and calibrated the parameters to target several moments in U.S. wage data. One concern is that our parametric assumptions have unreasonable implications relative to the entire distribution of wages in U.S. data. Figure 4 plots the empirical probability distribution function from the U.S. wage data (dashed line) and from data generated by

[18]In Appendix C, we illustrate this in more detail.

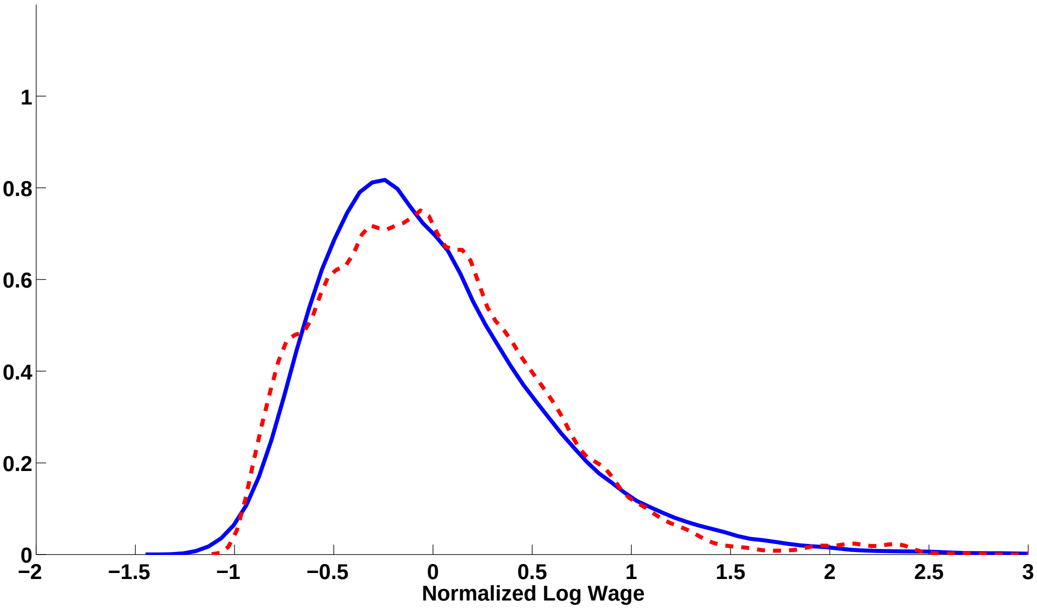

Figure 4: Aggregate Distribution of Wages in Data and Model

our model (solid line). They track each other closely, particularly with regards to the lower and upper tails and the general shape. In contrast, a model with log-normally distributed individual productivities features lower and upper tails that are more compressed in the model than in the data. This parallels the findings of Heckman and Sedlacek (1985), who also conclude that a log-normal model performs poorly in matching the tails of the U.S. wage distribution.[19]

5.7. Implications for U.S. Structural Transformation

Our model has implications for structural transformations—i.e., large reallocations of labor from agriculture to non-agriculture as a country develops. We find that our model's relationship among GDP per worker, the share of labor in agriculture, and the relative price of agriculture is quantitatively consistent with those in U.S. times series data from the *Historical Statistics of the United States*.[20]

In 1880, 44 percent of labor was employed in agriculture in the U.S., and its GDP per worker was 1/10th the U.S. level in 2000. As the U.S. developed, the share of labor in agriculture declined systematically, down to two percent in 2000. We varied A to replicate levels of GDP

[19] Appendix C provides a more complete discussion, including the distribution of wages by sector and the predictions of the model with log-normally distributed individual productivities.

[20] Appendix B outlines the exact series used in this analysis.

per worker relative to 2000 (our baseline) and computed the share of labor in agriculture. When the model replicates the same relative income level between 2000 and 1880, the share of labor in agriculture is 31 percent versus the 44 percent in the data. Over the entire time period, our model generates a systematic decline in the labor share similar to that in the data.

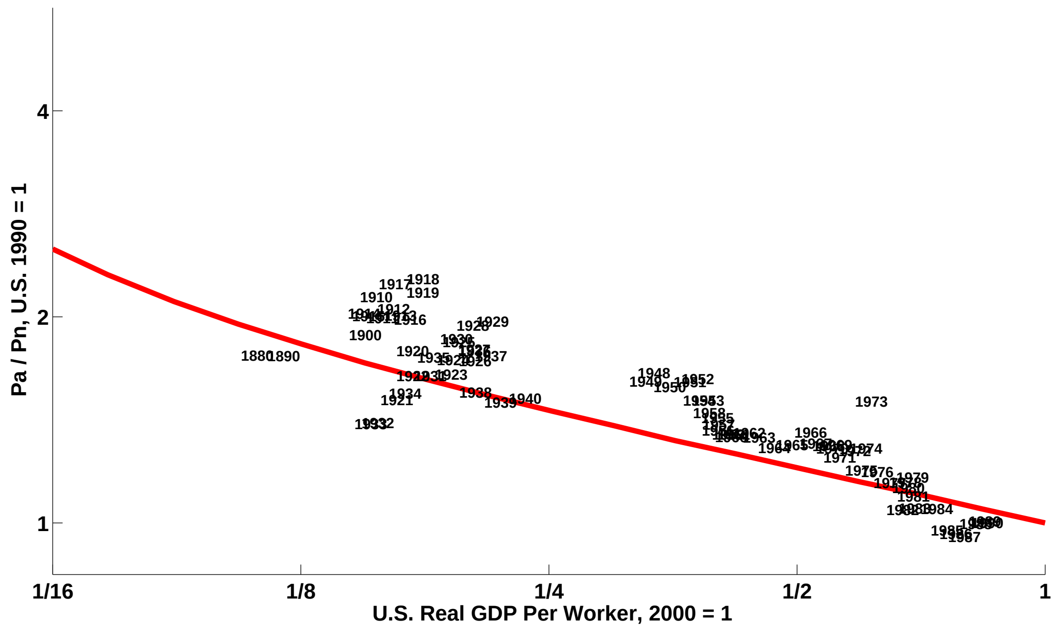

Figure 5: Relative Price of Agricultural Goods, U.S. Historical Data and Model

Figure 5 plots relative agriculture prices versus real GDP per worker for the U.S. Between 1880 and 1900 relative agriculture prices were around twice as high as they were in 1990. The solid line plots the implications of our model when we vary A to replicate levels of GDP per worker relative to 2000 (our baseline). All in all, the model's implications for relative agriculture prices over this time period line up reasonably well with the U.S. data.

6. Robustness of Quantitative Results

In this section, we ask how robust our quantitative results are to alternative modeling choices. We show that plausible alternative values for the parameters θ_a, θ_n, and ρ do not change the substantive nature of our conclusions. We then argue that the nature of out results would be unlikely to change if we relaxed the assumption that individual productivity has the same distribution in every country.

6.1. Individual Productivity Variance

One lesson of the illustrative example in Section 3 was that lower variance in productivity across individuals would lead to less predictive power for our model. We examine the sensitivity of our results once we re-calibrate our benchmark model to a plausibly smaller individual productivity variance. Recall that θ_a and θ_n were calibrated so that the model matches the variance in wages within each sector in the U.S. data. 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.

This observation motivates the following exercise. We follow the Postel-Vinay and Robin (2002) estimate, in an effort to be as conservative as possible, and recalibrate our model to match half the observed variance of log wages in agriculture and in non-agriculture. This implies that the standard deviation of the log wages due to individual productivity differences is 0.32 in agriculture and 0.40 in non-agriculture. In an effort to isolate the effect of lower individual productivity variance on our results, we recalibrate ρ to maintain the same implied linear correlation coefficient as in the baseline calibration, keeping the degree of dependence the same. We re-calibrate the preference parameters as described in Section 5.2.

Table 4: Labor Productivity Differences, Half-Variance Calibration

Ratio of 90th-10th Percentile			
Sector	Data	Model	Percent Explained
Agriculture	45	29	31
Aggregate	22	22	-
Non-Agriculture	4	11	61

Table 4 presents the results for agriculture and non-agriculture productivity differences for countries in the 90th and 10th percentiles of the income distribution. The model predicts that agriculture output per worker differences should be a factor of 29 in agriculture and 11 in non-agriculture. The third column of the table shows that this corresponds to the model explaining 31 percent of agricultural differences and 61 percent of non-agricultural differences, relative to aggregate differences. The key result from this exercise is that our mechanism still explains a sizeable fraction of the difference between aggregate and sector productivity gaps even when our individual productivity distributions have half the variance, as before.

6.2. Dependence of Individual Productivity

How does dependence of sector individual productivity affect our results? We compute the model's predictive power under a range of correlation coefficients running from 0.30 to 0.50 (compared to the 0.39 correlation in the benchmark model) by varying the dependence parameter ρ. In each case, we re-calibrate θ_n and θ_a to be consistent with the standard deviation of log wages in each sector and re-calibrate the preference parameters as described in Section 5.2. We do not attempt to match the ratio of average wages (since by varying ρ, we are no longer able to) but instead report the model's prediction for the sector wage ratio for each correlation value.

Table 5 shows the results of varying the correlation in individual productivity on the percent of difference between aggregate and sector output per worker differences. The first result of note in Table 5 is that the percent explained in non-agriculture is largely unaffected by changes in individual productivity dependence. The second noteworthy result is that decreasing the correlation increases the explanatory power in agriculture. For example, going from the baseline of 0.39 to a correlation of 0.30 increases the explanatory power from 54 percent to 72 percent. Increasing the correlation from 0.39 to 0.50 decreases the agriculture explanatory power from 54 percent to 29 percent. However, this also reduces the ratio of average sector wages down to the counterfactual level of 0.58, relative to 0.68 in the data.

Table 5: The Role of Dependence

Correlation in individual productivity	0.30	0.35	0.39*	0.45	0.50
Ratio of average wage $\bar{w}_a/\bar{w}_n$	0.72	0.70	0.68*	0.63	0.58
Percent Y/L explained in agriculture	72	62	54	37	29
Percent Y/L explained in non-agriculture	75	76	77	77	74

The intuition for this result is as follows. When ρ is low, agents with a high draw in one sector are not likely to have a high draw in the other sector. Thus, agents will be likely to have an absolute advantage in the sector in which they have a comparative advantage. As a result, lowering economy-wide efficiency lowers average individual productivity in agriculture—as agents with weaker comparative advantage in agriculture move into agriculture—and raises average individual productivity in non-agriculture. This outcome is consistent with the model satisfying both conditions of Proposition 2.

When ρ is larger, agents with a high draw in one sector are likely to also have a high draw in the other sector. Such agents are likely to have an absolute advantage in both sectors. Since the variance of individual productivity is higher in non-agriculture than in agriculture, these agents are more likely to have a comparative advantage at non-agriculture. Thus, lowering economy-wide efficiency raises average individual productivity in non-agriculture, as agents with weaker comparative and absolute individual productivity in non-agriculture move to agriculture. But the effects are less strong in agriculture, where many agents with a comparative advantage in agriculture have an absolute disadvantage in non-agriculture.

6.3. Individual Productivity Distribution Across Countries

In our baseline model, we assume that the distribution of individual productivity is the same across countries. We argue here that the substantive nature of results does not depend on this assumption. On the contrary, we make this assumption to emphasize that our results are driven by the allocation of productive individual, across sectors, not cross-country differences in individual productivity itself.

First consider relaxing the assumption that the mean of individual productivity in each sector is the same across countries. This could be the case if average years of schooling is lower in certain countries, as is well documented (e.g. Hall and Jones (1999)). This would be the same in our model as a lower value of A in the countries with lower average schooling. If average ability were systemically lower in agriculture than in non-agriculture in poor countries, this would lead to even lower agriculture output per worker, in addition to effects from our mechanism. One concrete example would be lower average schooling in agriculture, as studied by Chanda and Dalgaard (2008) and Vollrath (2009), and, hence, we abstract from it in our study.

Second, consider relaxing the assumption of identical individual productivity variance across countries. If the variance in individual productivity were systemically related to the level of development, then our results might overestimate or underestimate the quantitative importance of our mechanism. For example, if poorer countries had systemically less variance in individual productivity in both sectors, then our baseline model would tend to overstate our theory's importance. Nevertheless, using labor-income variance as a proxy for productivity variance, there does not seem to be a systematic correlation between GDP per capita and income variance. We conclude that our simplifying assumption of identical variance is not likely to systemically overstate our results.

7. Evidence: Women and Children in Agriculture

In this section, we provide direct evidence in support of our theory, using data on the composition of agriculture workers by sex and age across countries and over time. Under the assumption that women and children are at a comparative disadvantage in agriculture work compared to adult men, our theory implies that the fraction of agriculture workers that are women and children should be systematically higher in less-developed countries.[21] The reason is that, because women and children have this comparative disadvantage, they will work in agriculture only if subsistence food needs are strong enough.

7.1. Across Countries

To test this implication across countries, we use two data sources. First, we use data from the FAO on the composition of agriculture workers by sex in 162 countries. Second, we use census microdata from Integrated Public Use Microdata Series (IPUMS) to compute the composition of agriculture workers by sex and age for 38 countries. Using the FAO data, we compute the fraction of each country's agriculture workers that are female, and using the IPUMS data, we measure the fraction of agriculture workers that are either women or children (defined to be under 16 and economically active in agriculture.)

Figure 6 shows our calculations using the FAO data. We find that countries with higher shares of workers in agriculture tend to have a higher fraction of agriculture workers that are women, as our theory would predict. For the countries with 70 percent or more of workers in agriculture, roughly half of agriculture workers are women. In contrast, in countries with ten percent of workers in agriculture, on average, 30 percent of agriculture workers are women. A linear regression of the share of agriculture workers that are women on the share of workers in agriculture yields a precisely estimated slope coefficient of 0.29. Random assignment of workers or no difference in comparative advantage between men and women would imply a slope coefficient of zero; thus, Figure 6 provides evidence in support of our model.

The IPUMS evidence for women and children in agriculture reinforces this relationship. We find that the share of women and children systematically increases as the share of all

[21] Goldin and Sokoloff (1982, 1984) argue that women were at a comparative disadvantage in agriculture work in the U.S. in the 19th century. They argue that agriculture work is mostly strength-intensive, and hence favored men, while women and children were better suited to factory work, in which dexterity and small stature were valued traits. Foster and Rosenzweig (1996) provide complementary evidence from a sample of farmers from the Philippines, estimating men to be more productive than women in two types of agricultural tasks: plowing and weeding.

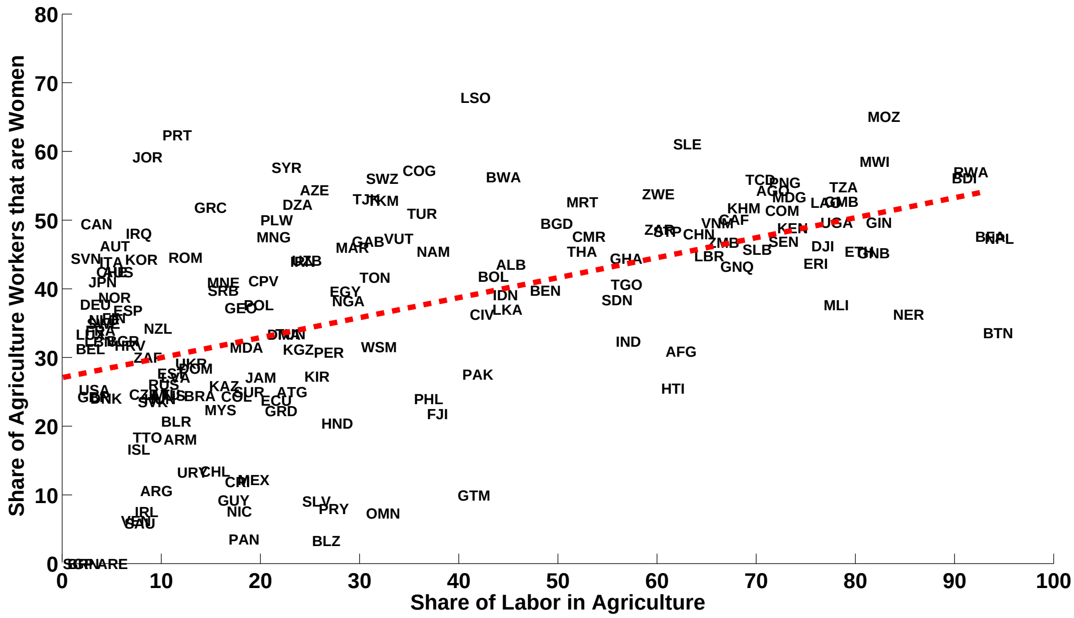

Figure 6: Share of Agriculture Workers that are Women

workers in agriculture increases. A linear regression of the share of workers that are female and children on the percent of workers in agriculture yields a precisely estimated slope coefficient of 0.38. This estimate implies that when a country moves from 70 percent to 30 percent of its workers in agriculture, the share of those workers that are female and children should decline from 50 percent to 34 percent. Again, this observation provides evidence in support of our model.

7.2. Across Time

Our model has implications for the composition of women and children in agriculture work during a structural transformation. In particular, it implies that when a structural transformation occurs (and labor moves out of agriculture), women should move out of agriculture at a faster rate than men. This implies that the share of women in agriculture should decline over the period during which the structural transformation occurs.

Evidence from Britain supports this prediction. Table 6 shows Allen's (1994) calculations for the composition of agriculture workers by age and sex in England and Wales during England's structural transformation, during which the share of workers in agriculture fell sharply. In 1700, 62.0 percent of agriculture workers in England were women and children and the balance adult men. This fell to 55.3 percent and to 36.3 percent by 1851. For women alone, the same figures were 32.5 percent in 1700, 30.3 percent in 1800, and 26.8 percent in

Table 6: Composition of English Agriculture Workers in Percent

	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)

1851. Men, on the other hand, went from representing just over one third of agriculture workers to just under two thirds.[22]

8. Extension with Land

In this section, we argue that the fixed supply of productive land in a country is an additional factor explaining why economy-wide efficiency differences would lead to even larger agriculture productivity differences across countries. The reason is that when a country is faced with low economy-wide efficiency and subsistence consumption needs, it must use some of its least-productive land for agricultural production. As efficiency rises, it can restrict its production to more and more productive plots of land. Thus, if productive land is a fixed factor of production in agriculture, economy-wide efficiency differences lead to larger output per worker differences in agriculture.

To formalize this argument, assume that all workers are identical, and let the production functions be

$$Y_a = AN_a^{1-\alpha}X^{\alpha} \quad \text{and} \quad Y_n = AN_n \tag{11}$$

where N_a and N_n are the number of workers in each sector and X is the quantity of land in the agricultural sector. Let α represent the extent of decreasing returns to scale in agriculture, where $0 < \alpha < 1$. For simplicity normalize the quantity of land to be 1. Let the 90th and 10th

[22]For the United States, Goldin and Sokoloff (1982, 1984) argue that sector differences in male and female individual productivities were a key factor in the industrialization of the northern U.S. in the 19th century. Their argument is that agriculture work in the North, which was mostly wheat and grain farming, was more strength-intensive than farm work in the South, where cotton and tobacco prevailed. In the South, dexterity and small stature were major determinants of a farm worker's ability, along with strength. Thus, women and children were at a much greater disadvantage compared to men in the North than the South. This was reflected in a smaller gender gap in agriculture wages in the South than in the North. As the United States grew, women became a much more important part of manufacturing in the North than the South. Goldin and Sokoloff (1982) find that, among the top 25 most female intensive manufacturing industries, 11 firms in the North for every one firm in the South, with a ratio of just 4.5 for other industries.

percentiles of countries have efficiency A^R and A^P. Then,

$$\frac{Y_a^R/N_a^R}{Y_a^P/N_a^P} = \frac{A^R}{A^P}\left(\frac{N_a^P}{N_a^R}\right)^\alpha . \tag{12}$$

Since α is less than one, and since countries with lower A have a higher share of workers in agriculture, productivity differences in agriculture are larger than the original A differences. The differences are bigger the closer α is to one.

While this amplification through land is qualitatively clear, we argue that its quantitative importance is likely to be modest. In short, the reason is that the importance of land in the agricultural production function is not large enough to generate both larger differences in agriculture productivity and much smaller differences in non-agriculture productivity.

The key parameter driving the extent of the amplification is α, land's importance in production. Using sector-level data, Valentinyi and Herrendorf (2008) compute the land share of production in agriculture to be 0.18. Using this value for α, the 90th- and 10th- percentile countries have three percent and 78 percent of workers in agriculture, which implies that $\left(\frac{N_a^P}{N_a^R}\right)^\alpha = 1.8$. This low degree of amplification implies that, when faced with efficiency differences of a factor of 22 across countries, agriculture will have output per worker differences of a factor of 39, which is roughly in line with the data. However, non-agricultural productivity differences are exactly a factor of 22, which is strongly counterfactual compared to the factor of four observed in the data. Thus, if this model can be broadly consistent with the agriculture differences in productivity, it does so at the expense of missing wildly on the rest of the economy.

An alternative is to consider exogenous efficiency productivity differences of a factor of four. In this case, the non-agriculture differences will also have a difference of a factor of four, which is consistent with the data. On the other hand, the agriculture differences will be a factor of 7.2 (4×1.8). Far lower than the factor of 44 in the data.

Nevertheless, we conjecture that adding land into our main model will lead to even more amplification of economy-wide efficiency differences into larger agriculture productivity differences. Intuitively, if both the quantity of high-productivity land and high-productivity people in agriculture is fixed, lowering economy-wide efficiency will lead an economy to use less productive types of both inputs, leading to lower measured output per agriculture worker.

9. International Trade

How will the model's predictions change if we allow for international trade? In this section, we argue that as long as a model with international trade generates labor allocations consistent with cross-country data, the model's quantitative predictions for sector productivity differences across countries will remain the same. The reason is that our results rely on the country differences in the allocation of workers across sectors, not necessarily on how that allocation arises in equilibrium.

This argument is clearly seen in the special case of our model described in Section 4.3. In this special case, sector productivity differences in agriculture are given by

$$\frac{Y_a^R/N_a^R}{Y_a^P/N_a^P} = \left(\frac{\pi_a^P}{\pi_a^R}\right)^{\frac{1}{\theta}} \left(\frac{A^R}{A^P}\right), \tag{13}$$

where the ratio of the labor allocations in agriculture determines the productivity difference across the two countries. Equation (13) says that if an open-economy model supports the same allocation of workers in agriculture as the closed-economy model, then the open-economy model's predictions for productivity differences are the same. The only distinction between the open- or closed-economy models is how they support the allocation of workers in equilibrium.

This suggests that any open-economy model must appeal to trade frictions to account for the difference in labor allocations across countries, while yielding plausible predictions for trade flows.[23] Results from the trade literature on the size of the trade frictions necessary to account for trade flows suggests that a closed-economy model is a reasonable approximation of a model with frictional trade (see Anderson and van Wincoop (2004) for a survey and Waugh (2009) for a focus on trade between rich and poor countries).

We want to emphasize, however, that understanding what prevents poor countries from trading themselves out of their plight is still an interesting question. In fact, we view this as one of the key questions in development economics and are currently exploring this topic.

[23]Frictionless trade would introduce strongly counterfactual outcomes. First, the shares of labor in agriculture (and non-agriculture) would be equated across countries. Second, because the shares of labor are the same in both countries, the extent of specialization would be the same, and ,hence, sectoral labor productivity differences would be the same. Both are true because: (i) frictionless trade implies that all countries face the same relative price p_a^W, and (ii) the relative price p_a^W determines the allocation of workers across countries through the sector labor supply cutoff in equation (4). Therefore, the allocation of workers in each sector is identical, and sectoral labor productivity differences are identical.

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

Our theory has new implications for the way economists think about the importance of agriculture for understanding aggregate 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 economy-wide efficiency in the face of subsistence agriculture requirements. In this case, it is optimal to employ many workers in agriculture who are less able in farm labor than in 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, Johnson, and Robinson (2002, 2001)).

A. Model Appendix

1.1. 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 economy-wide 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 (4), 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 (5), 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', (5), is for more agents to supply labor in the non-agriculture sector in economy A^2 than economy A^1. By (4), this occurs if and only if $p_a^2 < p_a^1$. ■

1.2. Proof of Proposition 2

Assume that $E(z_a|z_a/z_n > x)$ is increasing in x. By (4) we know that for any agent i with individual productivities z_a^i and z_n^i, if i chooses to work in agriculture in country P then $z_a^i/z_n^i > 1/p_a^P$, and if i chooses to work in agriculture in country R then $z_a^i/z_n^i > 1/p_a^R$. By Proposition 1 we know that $p_a^P > p_a^R$. Hence, by our assumption, $E(z_a|z_a/z_n > 1/p_a^P) < E(z_a|z_a/z_n > 1/p_a^R)$. Thus

$$\frac{Y_a^R/N_a^R}{Y_a^P/N_a^P} = \frac{A^R}{A^P} \cdot \frac{E(z_a|z_a/z_n > 1/p_a^R)}{E(z_a|z_a/z_n > 1/p_a^P)} > \frac{A^R}{A^P}.$$

A similar result holds when $E(z_n|z_n/z_a > x)$ is increasing in x. ■

1.3. Deriving Analytical Results in Independent Fréchet individual productivities

The probability we want to derive is Prob$\{z_n \leq p_a z_a\}$. To do so, note that this probability productivity is represented by

$$\pi_a = \int_0^\infty \exp\{-\left(p_a z_a\right)^{-\theta}\} g(z_a) dza$$

where the first term in the integral is the cumulative distribution function for productivity in non-agriculture evaluated at random variable $p_a z_a$ and the second term $g(z_a)$ is the probindividual productivity distribution function for the productivity distribution in agriculture. The anti-derivative for this integral is given by:

$$\frac{1}{p_a^{-\theta}+1} \times \exp\{-(p_a^{-\theta}+1)z_a^{\theta}\}$$

With the anti-derivative and evaluating the integral yields

$$\pi_a = \frac{1}{p_a^{-\theta}+1}$$

and similar arguments yields

$$\pi_n = \frac{p_a^{-\theta}}{p_a^{-\theta}+1}$$

To compute the conditional average individual productivity in each sector, we make the following argument. First notice that the conditional probability productivity distribution for workers in non-agriculture is:

$$\text{Prob}\left\{z_n < z | z_n > p_a z_a\right\} = \frac{\text{Prob}\left\{z_n < z, z_n > p_a z_a\right\}}{\text{Prob}\left\{z_n > p_a z_a\right\}}$$

Then computing the probability productivities in the numerator and the denominator we have:

$$\frac{\text{Prob}\left\{z_n < z, z_n > p_a z_a\right\}}{\text{Prob}\left\{z_n > p_a z_a\right\}} = \exp\{-(p_a^{\theta}+1)z_n^{-\theta}\}$$

Notice that the conditional probability productivity distribution of workers in non-agriculture is itself Fréchet distributed with centering parameter $(p_a^{\theta}+1)$. Using this insight we can now compute the average individual productivity of non-agriculture workers conditional on working in non-agriculture:

$$E(z_a | p_a z_a > z_n) = (p_a^{\theta}+1)^{\frac{1}{\theta}}\gamma$$

where the constant γ is the gamma function evaluated at $\frac{\theta-1}{\theta}$. Similar arguments yield average individual productivity of agriculture workers conditional on working in agriculture:

$$E(z_a | p_a z_a > z_n) = (p_a^{-\theta}+1)^{\frac{1}{\theta}}\gamma$$

B. Data Sources

- **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 did not work at least 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 and Coleman (2001). 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.

C. Not for Publication: Additional Support for Quantitative Predictions of Model

3.1. Model's Food Price Predictions compared to Producer Price Data

We show that the model's prediction that food prices are higher in poor countries is also consistent with data on producer prices of food. While in principle producer prices are more directly comparable to the prices in our model, since producer prices do not include a distribution margin, in practice producer prices for agriculture and non-agriculture goods are available for a much smaller set of countries. Nevertheless, we find that relative producer prices of agriculture goods behave very similarly to relative consumer prices of agriculture goods.

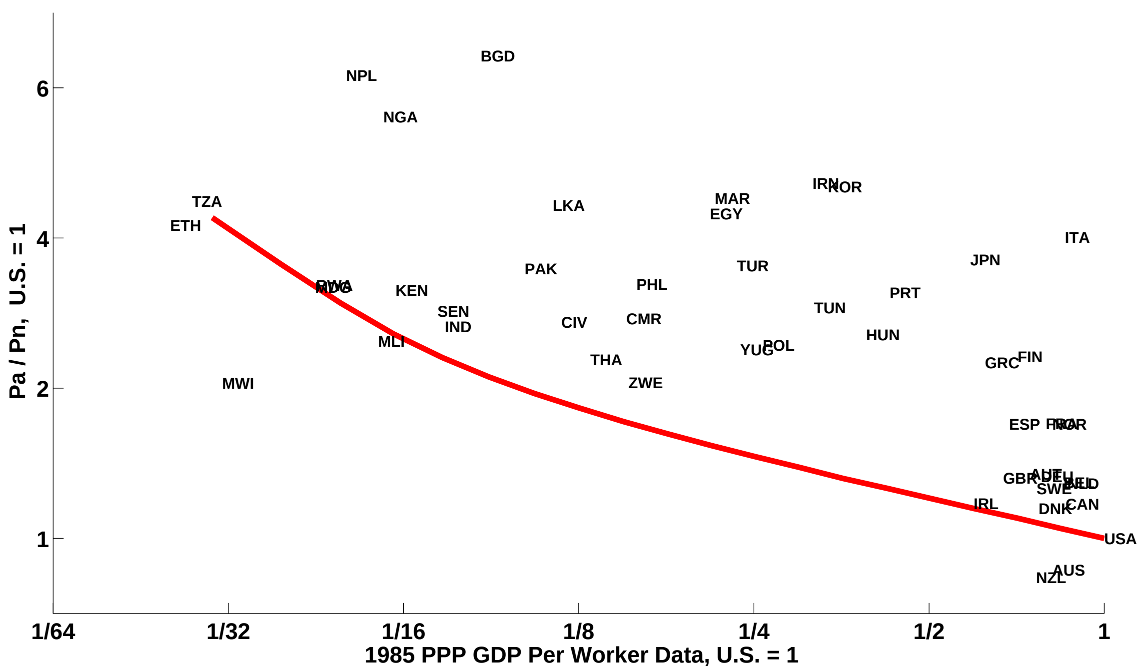

Figure 7: Relative Agriculture Producer Prices in Model and Data

Our data source is the 1985 FAO food producer price data, explored in detail by Adamopoulos (2009), and used by (Caselli (2005) and Restuccia, Yang, and Zhu (2008)) 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 in Figure 7. In the figure, one can see that relative prices of food is still higher in poor countries than rich countries, with the 10th percentile of countries around 4 times as high as in the United States (again normalized to one in the figure.) Note that relative agriculture prices appear a bit higher in poor countries once producer prices are used. This is consistent with the finding of Adamopoulos (2009) that distribution margins for food are moderately higher in richer countries than poor countries.

3.2. Log-Normal Version of the Model

In this section we show that the results presented in the text for our preferred specification of the model are similar to the predictions of a log normal version of our model. We show in addition that a log normal assumption on individual productivity, which has been used in some earlier studies, leads to predictions for the wage distributions that are worse than those of our preferred specification.

The log-normal version of the model assumes the individual productivity distribution is bivariate log-normal with variances σ_a^2 and σ_n^2 and covariance σ_{an}. As in the benchmark model we parameterize the three parameters using the variances of the sector wage distributions and ratio of average sector wages. The parameters which minimize the distance between these moments in the model and data are $\sigma_a^2 = 0.37$, $\sigma_n^2 = 0.30$, and $\sigma_{an} = 0.14$. The implied correlation of talent draws across individuals is 0.43, very much in line with the correlation from our preferred specification.

Table 7: Labor Productivity Differences, Log-Normal Model

Sector	Ratio of 90th-10th Percentile		
	Data	Model	Percent Explained
Agriculture	45	33.7	51
Aggregate	22	22	-
Non-Agriculture	4	14.5	42

Data Source: Caselli (2005)

Table 7 shows the results of lowering A in the log-normal model to match an aggregate GDP per worker difference of a factor 22. The model predicts a factor 33.7 difference in agriculture and a factor 14.5 difference in non-agriculture, which amounts to explaining 51 percent of the difference between the aggregate and the agriculture sector difference, and 42 percent of the aggregate and non-agriculture difference. While this is a bit less than the amount explained by the preferred specification of the model, it still represents sizeable predictive power. We conclude that the quantitative importance of the model is robust to assuming a log-normal individual productivity distribution.

Figures 8(a), 8(b) and 8(c) show the wage distributions in the calibrated log-normal model compared to the entire wage distribution in the data, and the distribution by sector. While in general the fit is good, the lower and upper tails are more compressed in the model than the data. This parallels the findings of Heckman and Sedlacek (1985), who also conclude that a log-normal model performs poorly in matching the tails of the U.S. wage distribution. Furthermore, the model has more weight on the lower half of the distribution than the data, particularly in agriculture.

For comparison, Figures 9(a), 9(b) and 9(c) show the wage distributions in our preferred specification of the model compared to the data. The fit of this model is better, particularly with regard to the tails. While the agriculture distribution still matches the data a bit less closely than the other two, all in all the preferred specification provides a close fit to the data.

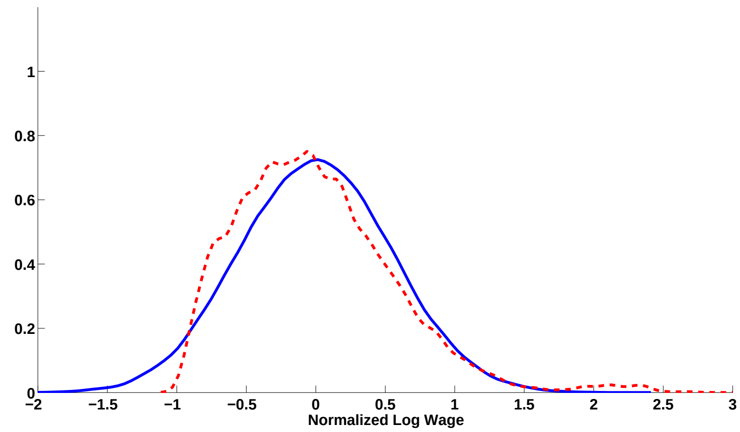

(a) Economy Wide Wage Distribution

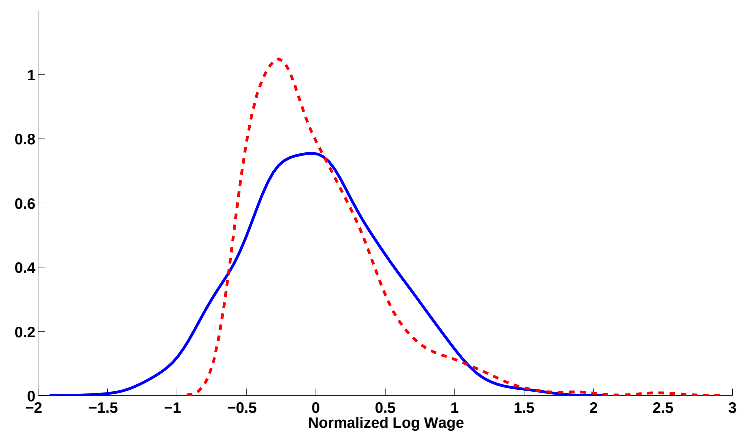

(b) Agriculture Worker Wage Distribution

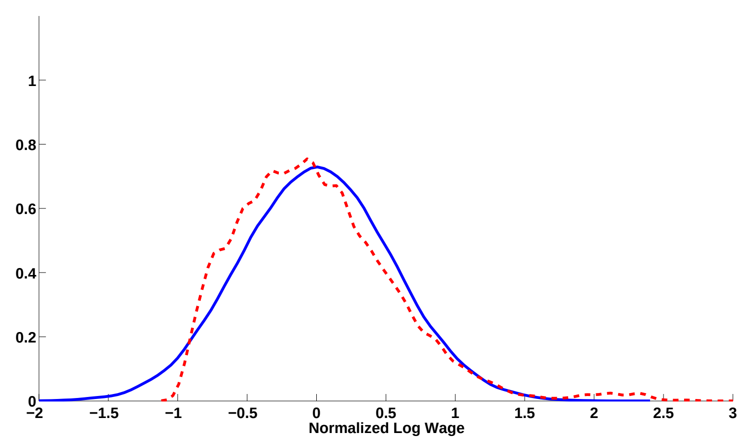

(c) Non-Agriculture Worker Wage Distribution

Figure 8: Wage Distributions from Log-Normal Model and Data

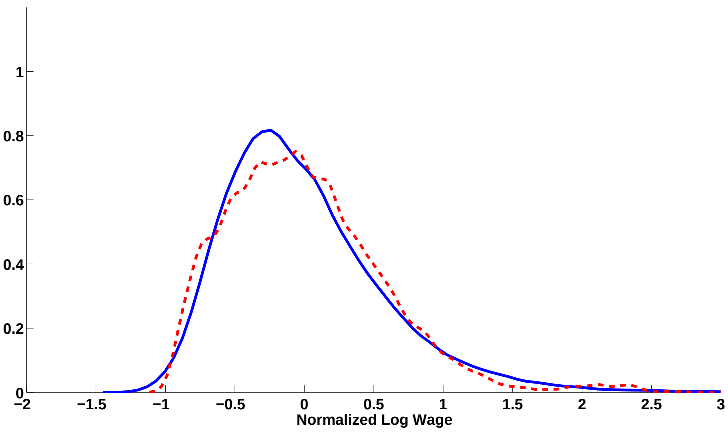

(a) Economy Wide Wage Distribution

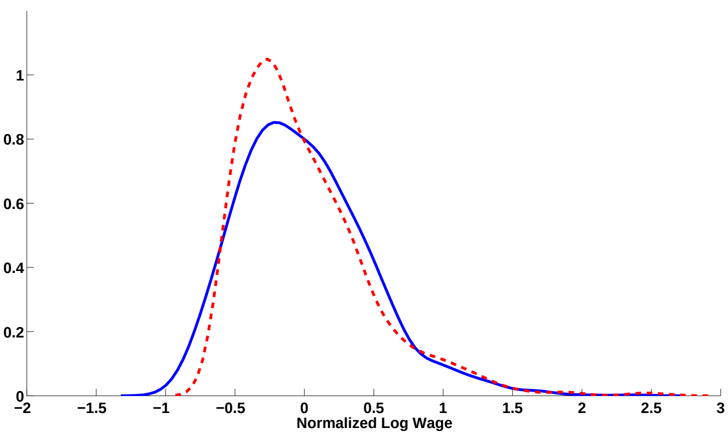

(b) Agriculture Worker Wage Distribution

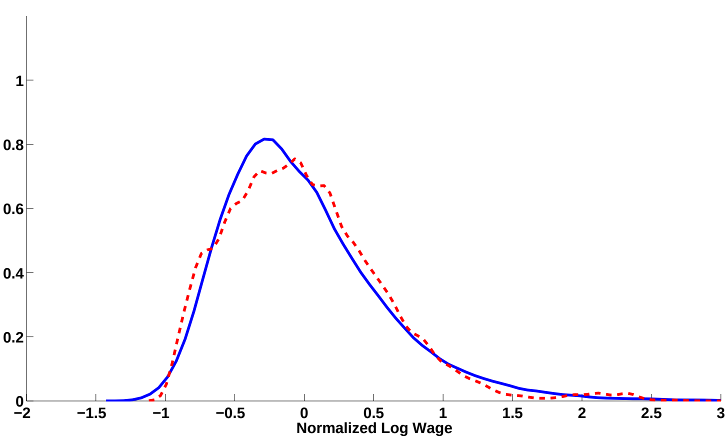

(c) Non-Agriculture Worker Wage Distribution

Figure 9: Wage Distributions for Preferred Model and Data

References

ACEMOGLU, D., S. JOHNSON, AND J. ROBINSON (2001): "The colonial origins of comparative development: An empirical investigation," *American Economic Review*, 91(5), 1369–1401.

ACEMOGLU, D., S. JOHNSON, AND J. ROBINSON (2002): "Reversal of Fortune: Geography and Institutions in the Making of the ModernWorld Income Distribution," *Quarterly Journal of Economics*, 117(4), 1231–94.

ADAMOPOULOS, T. (2009): "Do Richer Countries Have Higher Distribution Margins?," Unpublished Manuscript, York University.

ALLEN, R. (1994): "Agriculture During the Industrial Revolution," in *The Economic History of Britain Since 1700*, ed. by R. Floud, and D. McCloskey, vol. 1, chap. 5, pp. 96–122. Cambridge University Press.

ANDERSON, J., AND E. VAN WINCOOP (2004): "Trade Costs," *Journal of Economic Literature*, 42(3), 691–751.

ATKESON, A., AND M. OGAKI (1996): "Wealth-Vary Intertemporal Elasticities of Substitution: Evidence from Paneland Aggregate Data," *Journal of Monetary Economics*, 38, 507–534.

CASELLI, F. (2005): "Accounting for Cross-Country Income Differences," in *Handbook of Economic Growth*, ed. by P. Aghion, and S. Durlauf.

CASELLI, F., AND W. J. COLEMAN (2001): "The U.S. Structural Transformation and Regional Convergence: A Reinterpretation," *Journal of Political Economy*, 109(3), 584–616.

CHANDA, A., AND C.-J. DALGAARD (2008): "Dual Economies and International Total Factor Productivity Differences:Channelling the Impact from Institutions, Trade and Geography," *Economica*, 75, 629–661.

EATON, J., AND S. KORTUM (2002): "Technology, Geography and Trade," *Econometrica*, 70(5), 1741–1779.

FOSTER, A. D., AND M. R. ROSENZWEIG (1996): "Comparative Advantage, Information and the Allocation of Workers to Tasks:Evidence from Agricultural Labor Markets," *Review of Economic Studies*, 63, 347–374.

Goldin, C., and K. Sokoloff (1982): "Women, Children, and Industrialization in the Early Republic: Evidence fromthe 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, D., S. L. Parente, and R. Rogerson (2004): "Farm work, home work and international productivity differences," *Review of Economic Dynamics*, 7.

——— (2007): "The Food Problem and the Evolution of International Income Levels," *Journal of Monetary Economics*, 54, 1230–1255.

Graham, B. S., and J. R. Temple (2006): "Rich Nations, Poor Nations: How Much Can Multiple Equilibria Explain?," *Journal of Economic Growth*, 11(1).

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

Heathcote, J., F. Perri, and G. Violante (2009): "Unequal we Stand: An Empirical Analysis of Economic Inequality in the UnitedStates, 1967-2006," Prepared for *Review of Economic Dynamics* Special Issue.

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

Heckman, J. 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.

Herrendorf, B., and A. Teixeira (2009): "Barriers to entry and development," *Manuscript, Arizona State University and Fucape*.

Herrendorf, B., and A. Valentinyi (2009): "Which Sectors Make the Poor Countries So Unproductive?," Unpublished Manuscript, Arizona State University.

Hsieh, C.-T., and P. J. Klenow (2007): "Relative Prices and Relative Prosperity," *American Economic Review*, 97(3), 562–585.

Jones, C. I. (1994): "Economic growth and the relative price of capital," *Journal of Monetary Economics*, 34(3), 359–382.

KLENOW, P. J., AND A. RODRÍGUEZ-CLARE (1997): “The Neoclassical Revival in Growth Economics: Has it Gone Too Far?,” in *NBER Macroeconomics Annual 1997*, ed. by B. S. Bernanke, and J. Rotemberg. MIT Press, Cambridge, MA.

NELSEN, R. (2006): *An introduction to copulas*. Springer Verlag.

PARENTE, S. L., AND E. C. PRESCOTT (2000): *Barriers to Riches*. MIT Press, Cambridge.

POSTEL-VINAY, F., AND J.-M. ROBIN (2002): “Equilibrium Wage Dispersion with Worker and Employer Heterogeneity,” *Econometrica*, 70(6), 2295–2350.

RAO, P. (1993): “Intercountry comparisons of agricultural output and productivity,” *Etude FAO: Developpement Economique et Social (FAO)*.

RESTUCCIA, D., AND C. URRUTIA (2001): “Relative prices and investment rates,” *Journal of Monetary Economics*, 47(1), 93–121.

RESTUCCIA, D., D. T. YANG, AND X. ZHU (2008): “Agriculture and Aggregate Productivity: A Quantitative Cross-Country Analysis,” *Journal of Monetary Economics*, 55, 234–250.

ROSENWEIG, M. R., AND K. I. WOLPIN (1993): “Credit Market Constraints, Consumption Smoothing, and the Accumulation ofDurable Production Assets in Low-Income Countries: Investments in Bullocksin India,” *Journal of Political Economy*, 101, 223–244.

ROY, A. (1951): “Some Thoughts on the Distribution of Earnings,” *Oxford Economic Papers*, 3, 135–46.

SCHULTZ, T. (1953): *The economic organization of agriculture*. McGraw-Hill.

SUMMERS, R., AND A. HESTON (1991): “The Penn World Table (Mark 5): An Expanded Set of International Comparions,1950-1988,” *The Quarterly Journal of Economics*, 106.

VALENTINYI, A., AND B. HERRENDORF (2008): “Measuring factor income shares at the sectoral level,” *Review of Economic Dynamics*, 11, 820–835.

VOLLRATH, D. (2009): “How Important are Dual Economy Effects for Aggregate Productivity?,” *Journal of Development Economics*, pp. 325–334.

WAUGH, M. E. (2009): “International Trade and Income Differences,” *American Economic Review*, forthcoming.