Specialization, Economic Development and Aggregate Productivity Differences

David Lagakos

Federal Reserve Bank of Minneapolis,
and Arizona State University

Michael E. Waugh

Federal Reserve Bank of Minneapolis,
and New York University

July 2009

PRELIMINARY AND INCOMPLETE,
Comments Welcome

ABSTRACT

Cross-country labor productivity differences are large in agriculture and much smaller in non-agriculture. We argue that these relative productivity differences arise when subsistence consumption needs prevent workers in poor countries from specializing in the sector in which they are most productive. We formalize our theory by embedding the Roy (1951) model of ability into a two-sector general-equilibrium growth model in which the agents' preferences feature a subsistence food requirement. The model predicts that productivity differences in agriculture will be relatively larger than in non-agriculture, even though countries differ only in a sector-neutral efficiency term. A parameterized version of our model suggests that our theory is quantitatively important in explaining why agriculture productivity differences are so large relative to those in non-agriculture.

Email: lagakos@asu.edu, waugh@minneapolisfed.org. For helpful comments we thank Narayana Kocherlakota, Erzo Luttmer, Ellen McGrattan, Guillermo Ordoñez, Jim Schmitz and participants at the 2009 Midwest Macro Meetings. All errors are our own. The views expressed herein are those of the author and not necessarily those of the Federal Reserve Bank of Minneapolis or the Federal Reserve System.

1 Introduction

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

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

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

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

We then parameterize the model to study its quantitative implications. Our key assumption is that workers ability are drawn from independent Fréchet distributions with a common shape parameter. We show how this assumption implies that the share of labor in agriculture relative to non-agriculture is log linearly related the relative price of agriculture goods—consistent with the data. Furthermore, the relative labor share's elasticity with respect to relative prices identifies the shape parameter disciplining the degree of worker heterogeneity.

Our main exercise is to allow sector-neutral efficiency to vary across countries and study the impled differences in agriculture and non-agriculture productivity. When we feed sector-neutral efficiency differences into our model generating a factor of 22 difference in aggregate income, which corresponds to the differences between the 90th and 10th percentile of country income distribution, the model predicts a factor of 36.5 gap in agriculture labor productivity, and a factor of 13 gap in non-agriculture labor productivity. In the data there is a factor of 45 gap in agriculture and a factor of 4 gap in non-agriculture. Thus, our model explains sixty percent of the productivity differences in agriculture and nearly fifty percent of the productivity differences in non-agriculture between the 90th and 10th percentile of countries. We find that our model explains less of the sector productivity differences between the 90th percentile and countries with intermediate income levels, because differences in shares of labor in agriculture (and hence specialization differences) are smaller between these countries. We conclude that a completely frictionless economy with countries differing only by sector-neutral efficiency can generate meaningfully large differences in agriculture productivity and small differences in non-agriculture productivity between the richest and poorest countries.

We also show that the model performs quantitatively well in matching relevant development facts. Specifically, the model matches the cross-country data on the share of labor in agriculture, and it performs moderately well in matching agriculture's share in gross domestic product. We also compare our model's predictions relative to the U.S. structural transformation and show how it replicates the non-linear relationship between relative prices and the share of employment in agriculture. A failure of our model is its inability to generate a relative agriculture wage that decreases in income, as found in the data. We discuss how a simple modification of the model with schooling costs that decline with development, in the spirit of Caselli and Coleman (2001), can help improve our model's fit in this dimension.

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

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

2 Agriculture's Role in Accounting for Aggregate Productivity

In this section we highlight the important role of agricultural in understanding aggregate productivity. Specifically, we reproduce the findings of Caselli (2005) to illustrate how differences in labor productivity and shares of workers in agriculture account for much of the variation in aggregate output per worker across countries.

Table 1: Agriculture and Labor-Productivity Accounting

Panel A: Labor Productivity Differences

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

Panel B: Percent of Labor in Agriculture

Country Income Percentile	Percent
90th	2.8
10th	78.3

Source: Caselli (2005)

Panel A of Table 1 shows that labor productivity differences in agriculture are larger than aggregate differences, and that non-agriculture productivity differences are much smaller. The ratio of agricultural output per worker in the 90th and 10th percentiles of the income distribution is

45, compared to just 4 in non-agriculture. As a frame of reference, the ratio for the aggregates is 22. Panel B summarizes the well-known fact that poor countries have a much larger fraction of their work force in agriculture. A country whose per-capita income is in the 90th percentile has just 2.8% of its workers in agriculture, while the 10th percentile country has 78.3% of its workers in agriculture.

These two facts together suggest that labor productivity differences in the aggregate are almost completely accounted for by differences in agricultural productivity and shares of labor in agriculture. Caselli formalizes this argument by computing the hypothetical variance of cross-country aggregate output per worker assuming that agricultural productivity in all countries were equal to the U.S. level. His answer is just a factor of 1.6, down from the actual factor of 22! In other words, international labor productivity differences would be virtually eliminated. A similar experiment computes the hypothetical variance of aggregate output per worker assuming all countries had the U.S. share of workers in agriculture. This experiment yields a factor of 4.2 differences between the 90th and 10th percentile, which again is vastly lower than the 22 seen in the data.

One potential explanation for labor productivity differences in agriculture is physical capital per worker differences across countries. Caselli argues that labor productivity differences almost entirely represent total-factor productivity (TFP) differences. As he puts it, "the factor-only model explains virtually nothing of the observed per-capita income variance in agriculture: it's entirely a story of TFP differences, even more so than for aggregate GDP." (Caselli, 2005, page 49.) In independent work, Chanda and Dalgaard (2008) perform a similar set a counterfactual exercises using capital stock data from agriculture and non-agriculture, and conclude that around 85% of international TFP differences can be accounted for by TFP differences in agriculture relative to non-agriculture.

The results of this section suggest that cross-country differences in agriculture productivity and the share of workers in agriculture are central factors in accounting for aggregate productivity differences. It remains an unanswered question *why* the agriculture sector exhibits so much more variation in productivity across countries than the non-agriculture sector.

3 A Theory of Relative Agriculture Productivity

In this section we develop a model of productivity differences in agriculture relative to non-agriculture. The main result is that sector-neutral efficiency differences across countries lead to relatively larger productivity differences in agriculture and relatively smaller differences in the non-agriculture sector than the efficiency differences themselves. Proofs of all results of are available in Appendix A.

3.1 Households

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

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

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

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

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

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

3.2 Production

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

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

in agriculture and non-agriculture, where A captures sector-neutral efficiency, and $\tilde{L}_a$ and $\tilde{L}_n$ represent the total number of effective labor units employed in the two sectors.

Let Ω^a and Ω^n denote the sets of agents electing to work in agriculture and non-agriculture. The sector aggregate labor inputs $\tilde{L}_a$ and $\tilde{L}_n$ are defined as

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

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

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

3.3 Producer and Agent Optimization

Agents take as given prices and a wage schedule which maps talents into sector-specific wage offers. The problem for an agent is first to pick which sector to work in, and then to maximize (1) subject to (2).

Because of competition in production markets, the schedule of wages offered to a worker with talents z_a^i and z_n^i is equal to:

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

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

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

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

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

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

Due to the subsistence consumption constraints, agents consume relatively more food when their wage is lower.

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

3.4 Productivity Differences Larger in Agriculture than Non-agriculture

In this section we illustrate the main result of the paper: for countries that differ in sector-neutral efficiency levels, agriculture productivity differences are larger than the efficiency differences, and non-agriculture productivity differences are smaller. We first establish that relative food prices decline in the efficiency level, which is important for the other results to follow.

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

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

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

Proposition 2 *Consider two economies, rich and poor, with efficiency terms A^R and A^P such that $A^R > A^P$. Then differences in agricultural labor productivity are larger than the efficiency differences, which are larger than non-agriculture labor productivity differences:*

$$\frac{Y_a^R/L_a^R}{Y_a^P/L_a^P} > \frac{A^R}{A^P} > \frac{Y_n^R/L_n^R}{Y_n^P/L_n^P}.$$

Figure **??** illustrates the intuition behind Proposition 2. The curved grey frontiers represent the production possibilities frontiers (PPF) in the rich economy and poor economy, with the rich economy's PPF is simply shifted out by a factor A^R/A^P from that of the poor economy. The equilibria of the two economies are depicted as points on the frontier, with the poor country

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

choosing a relatively higher ratio of Y_a to Y_n, and the rich economy choosing a point with a lower ratio.

The fact that the PPFs are concave is key to understanding Proposition 2. The concavity of the PPFs arise because of the heterogeneity in worker ability. For points with a high ratio of Y_a to Y_n, such as the poor economy's equilibrium, agriculture productivity is low because the high relative price of food assures that many workers with relatively low ability in agriculture choose to work in agriculture. On the other side, the high food price assures that the average worker in non-agriculture has relatively high ability, since by (5) only those with high absolute ability in non-agriculture enter that sector. Hence, non-agriculture productivity is relatively high. In the rich economy, in contrast, the lower relative food price assures that the average ability is higher in agriculture, and lower in non-agriculture, than in the poor country.

Note that for Propositions 1 and 2 to hold, worker heterogeneity and non-homothetic preferences are necessary. If we relax either assumption then neither relative prices nor sector productivity vary with A. With homogenous agents the PPFs would be linear, and only a relative price equal to one would prevail in equilibrium. Hence labor productivity would be unchanged in each sector as one moves along the PPF. With homothetic preferences (but heterogenous agents) the PPFs would be concave, but the optimal production bundles of poor and rich countries would lie along one ray from the origin. Since shifting out the frontiers in sector-neutral way would lead to the same PPF slope along any given ray, it follows that the rich and poor countries would have identical relative prices, and identical sectoral productivity.

4 Quantitative Analysis of the Model

In this section we study the quantitative properties of our model relative to the data. First, we asses a key prediction of the model that relative prices decrease with GDP per worker. Then we parameterize the model and use the structure imposed by the model, relative price and labor data to identify the key parameter within our model.

With the calibrated model we then asses the quantitative importance of the mechanism explained above for understanding agriculture and non-agriculture labor productivity differences across countries. We also assess the model's cross-country predictions for shares of labor in agriculture, shares of GDP in agriculture, relative prices of agricultural goods, and wages of agriculture workers relative to non-agriculture workers.

4.1 Relative Prices Decrease With GDP Per Worker

Proposition 1 argues that the relative price of agriculture declines with level of development. As a first step in our analysis, we examine this relationship using food prices constructed using 2005 data available from the International Comparison Programme (see Appendix D for details). Figure 1 plots this relationship. The vertical axis contains the relative price of agricultural goods (expressed in log base 2) with the U.S. value normalized to one and the horizontal axis plots real gdp per worker relative to the U.S. also in log base 2 scale.

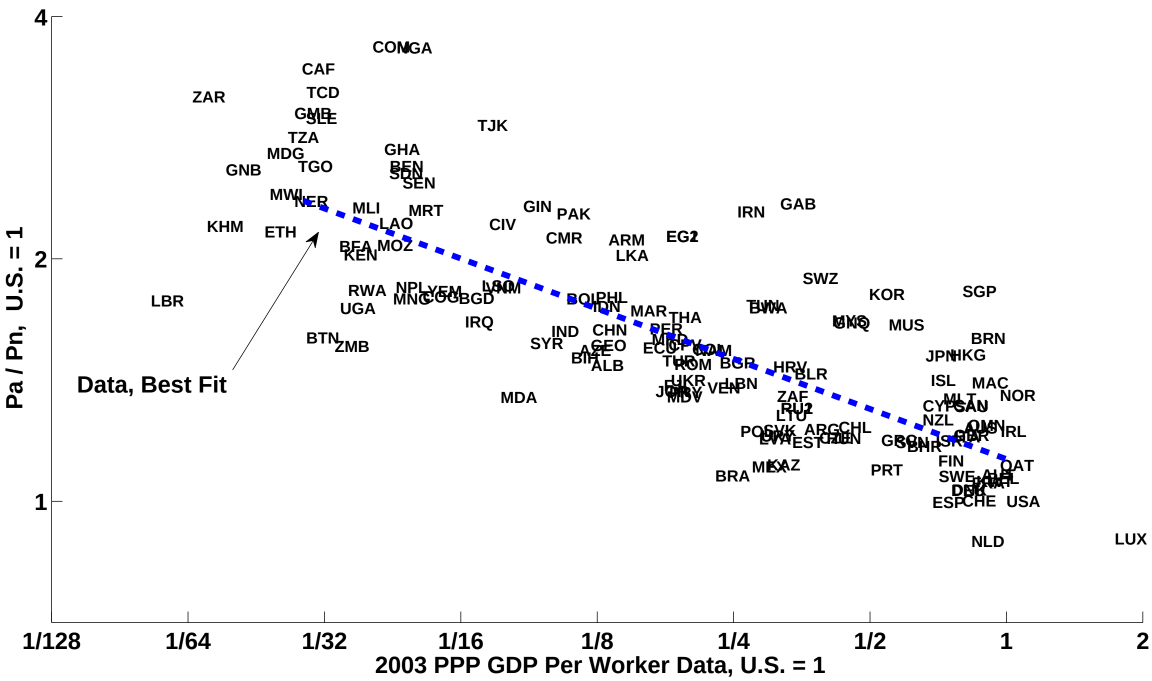

Figure 1: Relative Prices and GDP Per Worker

As predicted by Proposition 1, this feature is in fact captured by the data as well. The 10th percentile country in the data has a relative price of food that is roughly 2.5 times as high as the relative price in the United States, and the trendline shows a clear decreasing relative price in GDP per worker. Furthermore, not only is this relationship useful in confirming a property of our model, but as discussed in the next section we will parameterize the model such that the behavior of relative prices across countries will inform us on key parameters of the model.

4.2 Parameterization

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

We parameterize the model by assuming that z_a and z_n are drawn independently from a Type II extreme value distribution or Fréchet distribution:[2]

$$G(z_a) = e^{-\lambda_a z_a^{-\theta}} \quad \text{and} \quad G(z_N) = e^{-\lambda_n z_n^{-\theta}}. \tag{7}$$

In this model, relative λ's have no meaning and simply amount to a choice of units. Thus from here on we normalize $\lambda_A = \lambda_N = 1$.

For simplicity and identification purposes we assume that the talent distributions have a common θ. The parameter θ controls the dispersion of ability level and is central to our analysis. Mechanically, a smaller (larger) θ yields more (less) variation in ability levels.

Our key rationale for using this distributional assumption is how it conveniently relates relative employment shares between agriculture and non-agriculture the relative price of agriculture and the elasticity of relative employment shares to relative prices identifies the parameter θ. To see this, the distributional assumptions allow one to analytically express the probability and hence the mass of workers in agriculture and non-agriculture. One can show that the probability that a worker is in agriculture is:

$$\pi_a = \text{Prob}\left\{Az_n^i \leq p_A A z_a^i\right\} = \frac{p_a^\theta}{p_a^\theta + 1} \tag{8}$$

and the probability that a worker is in non-agriculture is

$$1 - \pi_a = \text{Prob}\left\{Az_n^i \geq p_a A z_a^i\right\} = \frac{1}{p_a^\theta + 1} \tag{9}$$

Note how Proposition 1 can be clearly seen in equation (9) — as the relative price of agriculture goods increases the share of workers in the non-agriculture sector decreases. Dividing the equation (8) by equation (9) and taking logs yields the following equation:

$$\log\left(\frac{L_a}{L_n}\right) = \theta \log(p_a) \tag{10}$$

Equation (10) says relative employment between agriculture and non-agriculture is related to the relative price with the elasticity θ.

The intuition behind this relationship is that with a low θ (meaning high productivity dispersion) some agents are much more productive in one sector relative to another. Thus the econ-

[2] This distribution has been used to parameterize Ricardian models of international trade originating with Eaton and Kortum (2002). One justification for this distribution is that it is an maximal extreme value distribution, thus representing an agents best activity within a sector. It is also very convenient. In trade models, the Fréchet distribution yields a log-linear gravity equation relating trade flows to structural parameters. Similarly in our framework, the Roy model plus our assumptions yield a log-linear equation relating employment shares, the relative price of agriculture goods, and structural parameters; see equation (10).

omy requires large changes in the relative price of agriculture to induce agents to switch sectors. If there is a high θ (meaning low productivity dispersion) agents are more similar in productivity across sectors and hence only small changes the relative price of agriculture are required to induce agents to switch sectors.

4.3 Estimating θ

The construction of employment shares and relative prices are discussed in the data appendix. As discussed above, our approach is to use the structural relationship described in Equation (10) in conjunction with data to estimate θ.

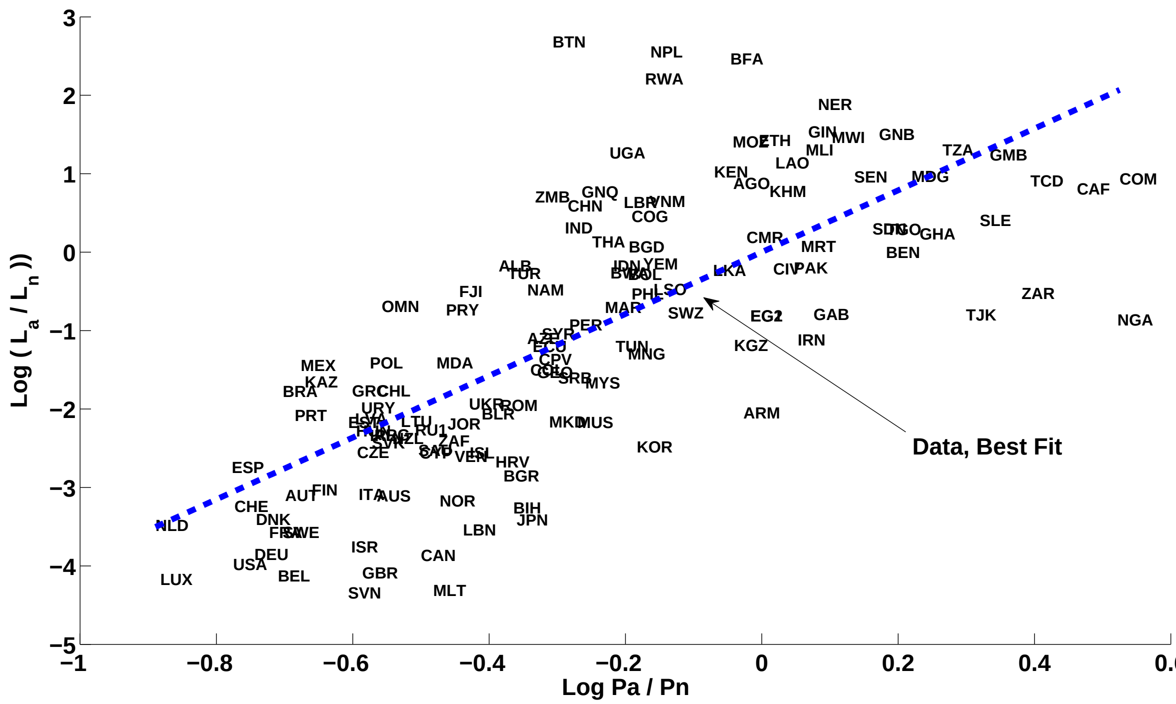

Figure 2: Relative Prices and Relative Labor Shares

Figure 2 plots log relative prices on the x-axis and log relative employment shares on the y-axis and the best fit line from the regression. Clearly there is a strong positive relationship between these two variables and other than several outliers the relationship appears to be linear in logs as predicted by the model. OLS yields an estimate of θ equal to 3.94 and is statistically different from from zero at the 1 percent level. This will be our baseline value throughout the rest of the analysis.[3]

[3] In the appendix we plot and estimate a similar relationship using only U.S. time series data giving a similar estimate of θ. Though in estimated in a very different context, Eaton and Kortum (2002) and Bernard, Eaton,

This estimate of θ has implications for observed the observed distribution of income within a country. One widely available measure for a cross-section of countries is the Gini coefficient. Our estimate of θ implies a Gini coefficient of 0.19. Bandourian, McDonald, and Turley (2002) find Gini coefficients for a wide range of countries between 0.20 and 0.50. While we are on the lower end of these estimates, this is result is reassuring for two reasons. First, many other factors may go into determining the distribution of income within a country other than productivity which could increase the dispersion in income. Second, because more dispersion in productivity will translate into lager amounts of productivity differences explained across countries, our estimate will generate results that are conservative. We are currently exploring the implications of our model for variability in wages and income with in a country more completely.

4.4 Preference Parameters

For the preference parameters, we pick ν to give a long-run food expenditure share of 2%. This choice is conservative relative to other similar models of structural change: Caselli and Coleman (2001) and Duarte and Restuccia (2008) pick a value of 1%, while Restuccia, Yang and Zhu (2008) pick a value of 0.5%. Admittedly, our model's results are sensitive to the choice of this value with higher (lower) values allowing us to explain less (more) of the variation in agriculture and non-agriculture labor productivity.

We set $\bar{a}$ to match a subsistence consumption need of 34% of average income in a model country with 7.5% of the U.S.'s per capita GDP. This is consistent with the independent estimates of subsistence food consumption requirements of Rosenzweig and Wolpin (1993), and Atkeson and Ogaki (1996), both of which use panel data from a sample of rural housholds from India (which had 7.5% of the U.S. per capita GDP in 1984).[4]

4.5 Quantitative Predictions for Sector Productivity Differences

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

Jensen, and Kortum (2003) find values of θ equal to 3.6 which is close to our value. A distinguishing feature of our analysis is the clean identification of θ while in trade models θ is often not well identified and various approaches are necessary to pin it down.

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

Table 2: Labor Productivity Differences

Sector	Ratio of 90th-10th Percentile		
	Data	Model	Percent Explained
Agriculture	45	36.5	63
Aggregate	22	22	-
Non-Agriculture	4	13.3	48

Data Source: Caselli (2005)

Table 2 shows the model's predictions for the ratio of the 90th to 10th percentile of countries in the model and data. The differences in aggregate output per worker (expressed as GDP per worker at Gheary-Khamis international prices) is a factor of 22 in the model and data by construction. The model predicts agriculture output per worker differences should be a factor of 36.5, and in non-agriculture it predicts a factor of 13.3 difference. In the data these ratios are (as described in Section 2) a factor of 45 and 4 respectively. The third column of the table shows that this corresponds to the model explaining 63% of agricultural differences and 48% of non-agricultural differences, relative to aggregate differences. The results in Table 2 show that differences in patterns of specialization are quantitatively important to understanding relative sector productivity differences between the richest and poorest countries.

Table 3 illustrates the model's predictions for developing countries with relatively higher average income. Specifically, it shows the model's prediction for the 90th-50th ratio and 90th-25th ratio. In the latter case aggregate productivity in the model and data differ by a factor of 9.4, again by construction. In the 90th-25th case, the model predicts a factor of 15.2 in agriculture and 8.4 in non-agriculture, compared to 31.1 and 2.7 in the data. The model predicts 27% and 15% of the agricultural and non-agricultural productivity differences, relative to the aggregate, as in the data, which is still large, but substantially lower than the 90th-10th percentile ratio.

In the 90th-50th case, the model fares much worse. The aggregate differences are chosen to be a factor of 3.1 as in the data. The model predicts differences in agriculture and non-agriculture of 3.8 and 3.04, compared to 11.1 and 1.9 in the data. This amounts to explaining just 8% and 5% of the sector differences, respectively.

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

Table 3: Labor Productivity Differences – Intermediate Income Levels

Sector	Ratio of 90th.-25th Percentile		Percent Explained
	Data	Model	
Agriculture	31.1	15.2	27
Aggregate	9.4	9.4	-
Non-Agriculture	2.7	8.4	15
	Ratio of 90th-50th Percentile		
	Data	Model	
Agriculture	11.1	3.8	8
Aggregate	3.1	3.1	-
Non-Agriculture	1.9	3.04	5

Data Source: Caselli (2005)

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

4.6 Assessing The Model's Other Quantitative Predictions

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

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

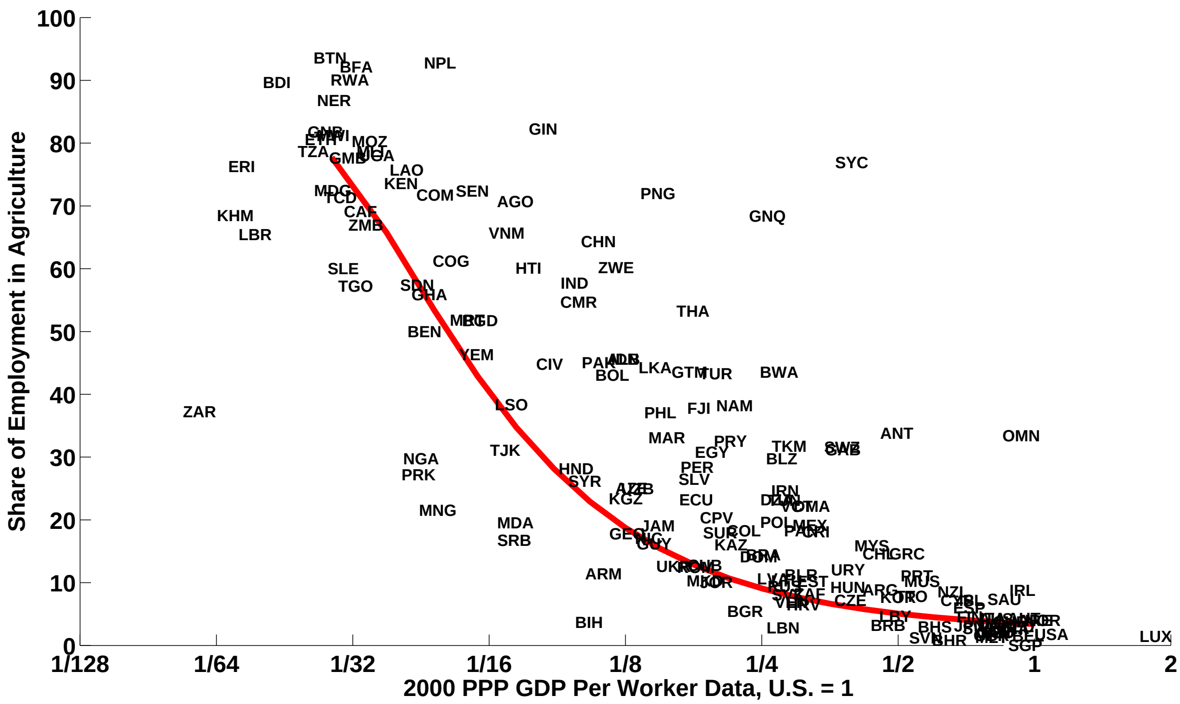

Figure 3: Shares of Labor in Agriculture in Model and Data

data is successfully captured by our model.

Next we turn, in Figure 4, to the model's predictions for the share of GDP in agriculture. While similar to the labor shares shown above, note that in the data the GDP shares in agriculture are systematically lower than the labor shares in agriculture. In Kenya, for example, agriculture employs 74% of the workers but produces just 28% of GDP. While our model does a reasonable job of capturing GDP shares for agriculture in the countries with around 1/8 the U.S. income level or higher, it substantially over-predicts the GDP share in agriculture in countries with lower income.

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

As another check of our model, we compared our models predictions relative to the U.S. time series of relative agriculture prices and labor shares in agriculture.[5] Figure 5 plots the relative price of agriculture on the vertical axis with the U.S. value in 1990 normalized to one. On the horizontal axis, the share of employment in agriculture is reported as well. The model's prediction is plotted in the solid line. Note that the model correctly picks up the decline in

[5]The data appendix describes the source of the data used.

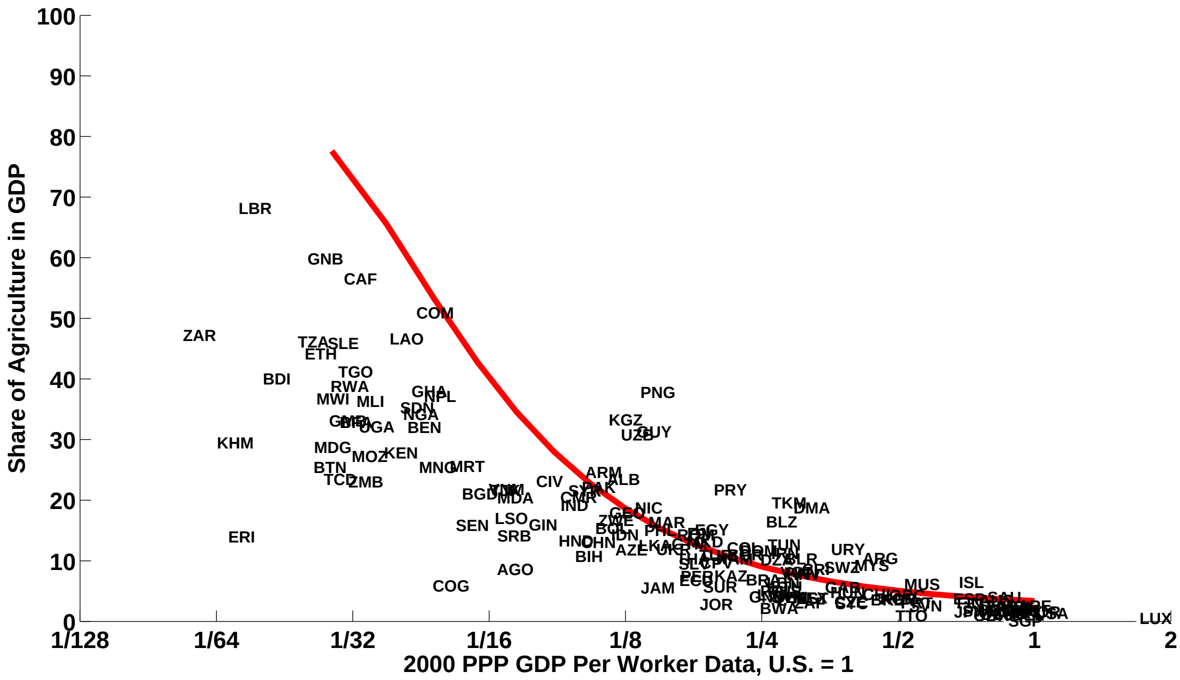

Figure 4: Agriculture Share in GDP in Model and Data

relative prices and the employment share in agriculture. More importantly our model replicates this fact in the non-linear way that the data exhibits.

The final prediction we consider is the average wage in agriculture relative to non-agriculture. This is the our model's key failure. Figure 6 shows the model's predicted relative agriculture wage along with data on the ratio of nominal average wages in agriculture compared to non-agriculture in a set of countries, which we collected from the International Labor Organization (ILO).

Two features are of note in the data. First, the average wage of agriculture workers is lower than in non-agriculture, with an average ratio of around 0.6. The second noteworthy aspect of the figure is that the relative agriculture wages increase in GDP per capita. In principle some of this upward slope could be since the wages are nominal, and there is evidence that the cost of living is lower in rural areas, particularly in less-developed countries, where transportation costs are highest. Our parametrization in contrast implies the relative wage of agriculture workers is the same in non-agriculture and that it is constant across level of development. This feature seems driven entirely by the Fréchet distribution over abilities. While this prediction of our model is somewhat mysterious, in the next section we explore potential ways to amend the model to be more in line with these data.

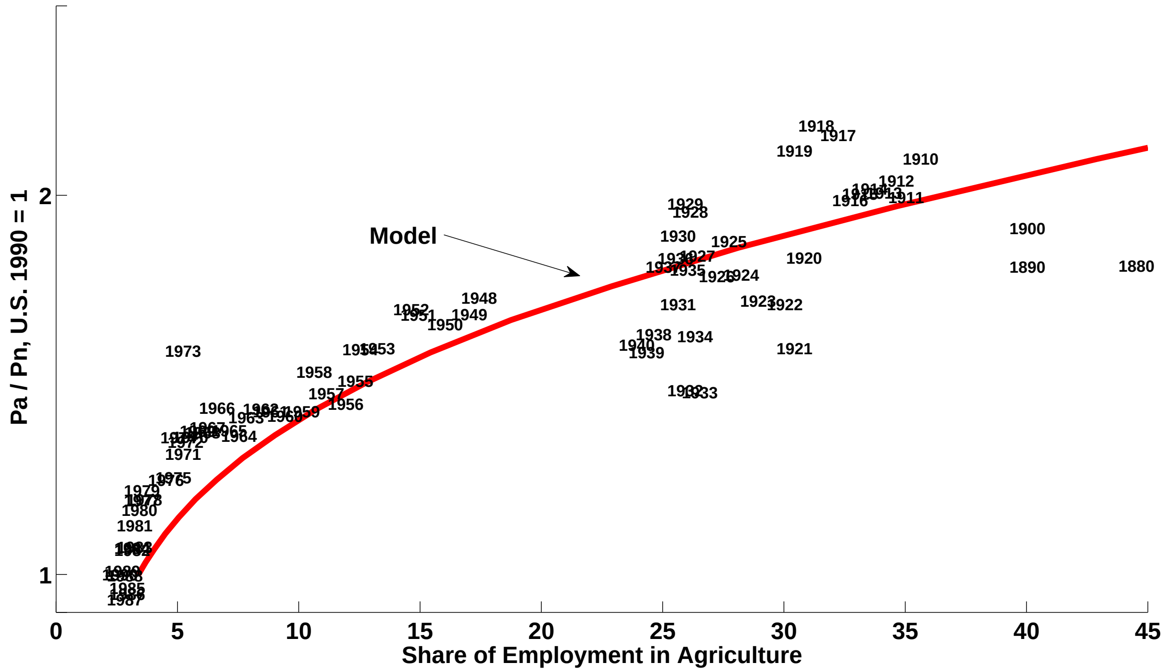

Figure 5: Labor Share in Agriculture and Relative Prices in Model and U.S. Historical Data

5 The Model's Limitations and Fixes

We first discuss how the model behaves once we add the possibility of trade, and second, we explore an extension which allows for the possibility of schooling as a way to improve the model's predictions for relative wages.

5.1 Adding Trade to the Model

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

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

Proposition 3 *Imagine that the rich and poor economies could buy or sell as much as they wanted on*

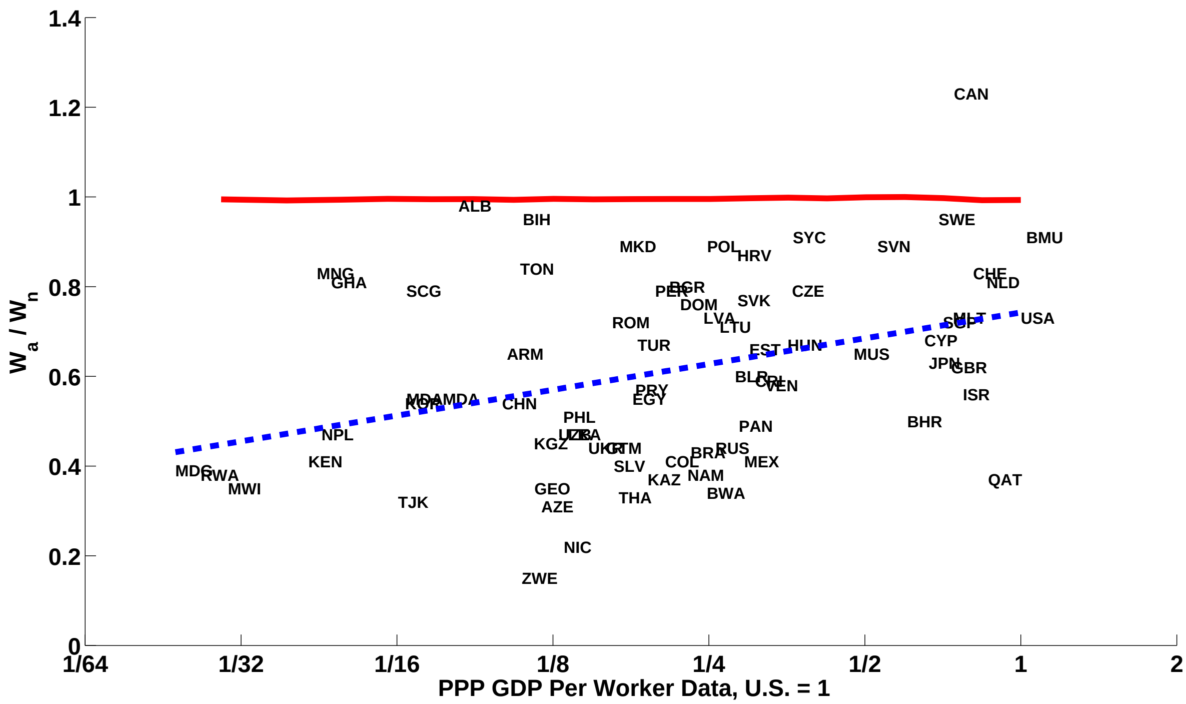

Figure 6: Average Wages in Agriculture Relative to Non-agriculture in Model and Data

world markets at a relative food price p_a^W. Then the following must hold:

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

Proposition 3 says that under frictionless trade, two things are true. First, the extent of specialization would be the same in both countries, and hence labor productivity differences between the rich and poor countries would be the same in agriculture and non-agriculture. Second, the shares of labor in agriculture would be equated across countries. Both are true because, under a common relative price, the sector labor supply cutoff (5) is identical in both countries, and hence the composition of workers in each sector are identical as well.

But the prediction of labor shares being equal across countries is strongly counterfactual. As is well known, a substantially higher fraction of labor in poor countries is in agriculture than rich countries. In Section 2, for example, we cite evidence that the United States has just 2.8% of its labor in agriculture compared to over 78.3% of the labor in a country at the 10th percentile of the world per-capita income distribution. Thus, we conclude that adding frictionless trade will generate counterfactual predictions.

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

5.2 Adding Schooling to the Model

Another potential way to improve the model's predictions, particularly for relative agriculture wages, would be to allow for the possibility of schooling. In the current model, as a country's efficiency increases, the sole reason labor moves from agriculture to non-agriculture is that the price of agriculture falls. Caselli and Coleman (2001) offer an additional factor that induces workers to leave agriculture, namely, a falling price of eduction, which increases the prospect of leaving agriculture to become a skilled non-agriculture worker. Caselli and Coleman argue that in the United States, the falling cost of education was a major factor leading to reallocation of workers into non-agricultural work. Furthermore, they show that without a falling education cost, standard models of structural change cannot reconcile the rising relative wages of non-agriculture workers which occurred over the late 19th and 20th centuries.

This finding suggests that allowing for the possibility of schooling in our model could help fix the model's counterfactual prediction regarding relative wages and level of development. Specifically, we could allow for schooling that increases a worker's effective number of labor units in non-agriculture work at some cost, where the cost of that schooling declines with general efficiency, A. Amending the model in this direction would obviously complicate our approach to identifying and change our estimate of θ. However, we view this as a potentially first-order extension of our model.

6 Historical Evidence: Males versus Females

In this section we provide some direct evidence in support of our theory. While many dimensions of ability heterogeneity are not observable, along two particular observable dimensions, namely age and sex, there is concrete evidence of ability differences in agricultural and non-agricultural tasks. Historians and development economists have argued that women and children generally have a comparative disadvantage at farm work than adult men. As one piece of evidence, Goldin and Sokoloff (1982, 1984) show that wages were much lower for women in farm work in the United States, earning roughly one third to one half as much as men in farm work in the nineteenth century. Foster and Rosenzweig (1996) provide complementary evidence from a sample of farmers from the Philippines, among which they estimate males to have an absolute productivity advantage in several types of agricultural tasks.[6]

Our theory thus predicts that women and children should have been the first to move off the farms and into non-agricultural work. This is in fact what happened. In Britain, according to Allen (1994), the fraction of farmers that were women or children declined substantially during Britain's industrial revolution. Table 4 shows Allen's calculations for the composition of farm workers in England and Wales between 1700 and 1851. In 1700, a full 62.0% of farm workers in England were women and children, with the balance adult men. By 1800 this percent fell to 55.3%, and by 1851 it was down to 36.3%.

Table 4: Composition of English Farm Workers

	1700	1800	1851
Men	38.3	44.7	63.7
Women and Children	62.0	55.3	36.3
Total	100.0	100.0	100.0

Data Source: Allen (1994)

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

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

[6]They estimate a one-factor model with two tasks: ploughing and weeding. They find that men had an absolute advantage in both and a comparative advantage in ploughing.

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

7 Effects of Policies That Restrict Labor Mobility

Our model implies that government policies that restrict migration within a country are likely to reduce productivity. One component of these policies that has received relatively less attention is the government's role in selecting *which* agents may move and which may not. Our theory says that letting the market allocate workers is more efficient than letting the government allocate them.

China is an important example of a growing country that is underdoing structural change and has restricted the flow of workers out of agricultural areas, typically in central China, into urban areas, typically on the eastern coast. One component of this *hukou* system is that the local and federal government officials decide which workers are permitted to migrate and which must remain in rural areas (See e.g. Au and Henderson (2006) and Lau, Qian, and Roldand (2000).)[7]

In future work we plan to use the model estimate the potential welfare cost of policies that restrict which individuals may leave agriculture and which may not. A rough idea is as follows. Imagine an economy like China in 1985 with 70% of its workforce in agriculture, and imagine it moves to a new steady state in which GDP per worker is (say) doubled. This corresponds roughly to Chinese growth from 1985 to the present. We can compute the welfare costs under two scenarios which restrict migration out of agriculture. The first randomly picks some fraction ω of agricultural agents that may move from agriculture to non-agriculture in such a way that markets clear, and the second lets the market reallocate agent. The experiment is in the same vein as that of Alder (2009), who assesses the efficiency losses from non-assortative matching of managers to projects.

[7] In addition to policies restricting which workers leave agriculture, there is evidence of social norms which restrict worker movements across sectors. Hayashi and Prescott (2008) cite evidence of social customs in the period before World War II which kept first born sons of farmers in agriculture.

8 Conclusion

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

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

A Model Appendix

A.1 Proof of Proposition 1

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

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

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

A.2 Proof of Proposition 2

Define $\rho(A)$ to be the relative productivity of agriculture to non-agriculture in an economy with efficiency A, i.e. $\rho(A) \equiv \frac{Y_a}{L_a} / \frac{Y_n}{L_n}$. It suffices to prove that $\rho'(A) > 0$. Recall that by (5), a agent with non-agriculture productivity z_n works in agriculture if and only if $z_a > z_n/p_a$, and that a agent with agriculture productivity z_a enters non-agriculture if and only if $z_n > z_a p_a$. Thus we can write

$$\rho(A) = \left(\frac{\int_0^\infty \int_{z_n/p_a}^\infty z_a \, g(z_a, z_n) \, dz_a \, dz_n}{\int_0^\infty \int_{z_n/p_a}^\infty g(z_a, z_n) \, dz_a \, dz_n} \right) \Big/ \left(\frac{\int_0^\infty \int_{z_a p_a}^\infty z_n \, g(z_a, z_n) \, dz_n \, dz_a}{\int_0^\infty \int_{z_a p_a}^\infty g(z_a, z_n) \, dz_n \, dz_a} \right).$$

By Proposition 1, we know that p_a is decreasing in A, and hence the left-side ratio is increasing in A, and the right-side ratio is decreasing in A. Thus $\rho'(A) > 0$. ■

B Estimating θ From U.S. Time Series

To bring in more evidence regarding this parameter, we also estimated θ using only U.S. data over a time series between 1880 and 1990. The data employed is described in the data appendix. Figure 7 plots data from the U.S. time series. On the horizontal axis are log relative prices and the vertical axis are relative labor shares. Note that the data appears to be log-linear in relative prices and relative labor shares. Using OLS, the estimate of θ is equal to 4.63 and is statistically different from zero at the one percent level. This estimate is slightly higher than that reported upon than using data from the cross-section of countries.

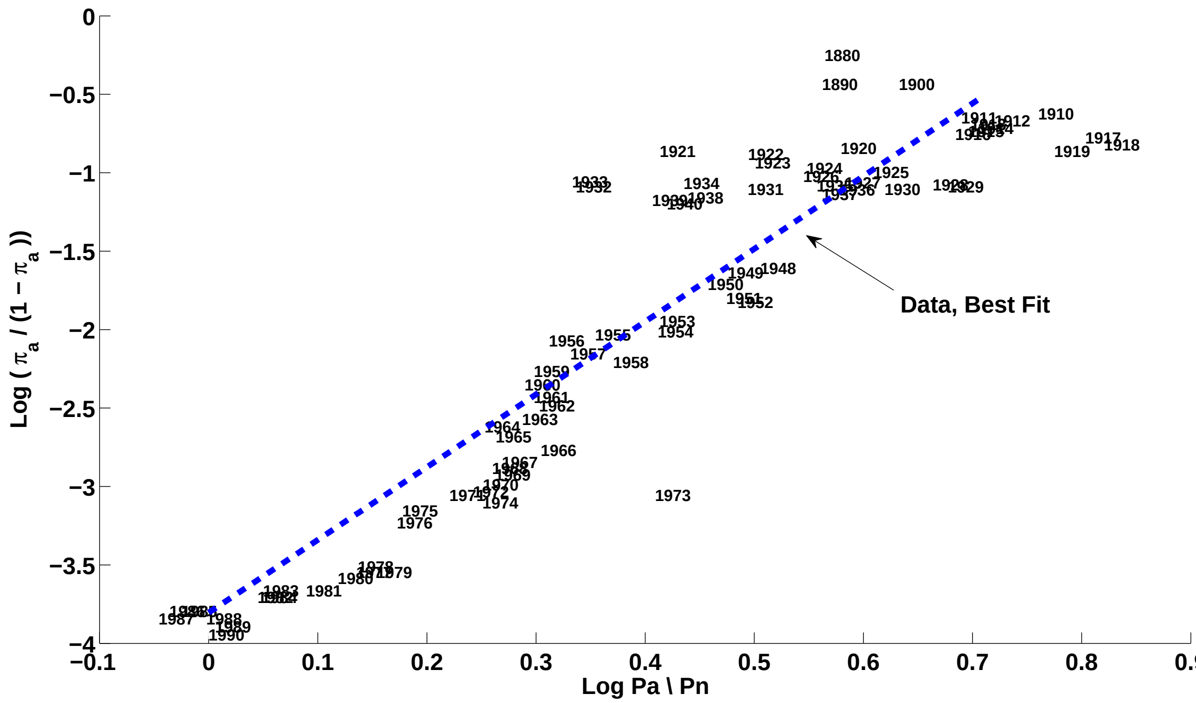

Figure 7: U.S. Time Series: Relative Prices and Relative Labor Shares

C The Failures of the Model with Homogenous Workers

In this section, we study the implications of shutting down the key innovation of our paper–heterogenous workers with occupational choice—and study a simple model with homogenous labor and allow for sector specific total factor productivity differences. With this simple and more standard model we study its implications with exogenous sectoral productivity differences and ask if it is reasonable. It is not. The reason for its failure is because large sectoral productivity differences imply implausibly large differences in relative prices and the model does not generate a log-linear relationship between relative labor shares and relative prices as exhibited in the data.

To illustrate this point, consider the simplification of our model with homogenous labor. Now the production functions are simply

$$Y_a = A_a A L_a \quad \text{and} \quad Y_n = A_n A L_n \tag{12}$$

with L_a and L_n representing the number of workers in each sector, not the number of effective labor units as in our model. Furthermore, now production functions have a total factor productivity term that is common A and a sector specific productivity term A_a and A_n. Preferences will remain the same as discussed in the text.

Two features of this model suggest it will have difficulties with respect to the data. The first difficulty with this model is that it dramatically over predicts the amount of variation in relative prices across countries. To see this, note that in this model prices are pinned down entirely by the production side with the relative price being $p_a = \frac{A_n}{A_a}$. Now compare the relative price between a country in the 10th percentile and 90th percentile of income per worker the relative price can be expressed as

$$\frac{p_a^{10}}{p_a^{90}} = \left(\frac{A_n^{10}}{A_n^{90}}\right) \times \left(\frac{A_a^{90}}{A_a^{10}}\right) \tag{13}$$

Plugging in numbers from Table 1 implies that the relative price between the 90th and 10th percentile should be a factor of $11 = .25 \times 45$. In the data the relative price between a country in the 90th and 10th percentile is a factor of 2.5 — significantly less than that predicted by the model.

The second difficulty is that the model does not generate a log-linear relationship between relative labor shares and relative prices as exhibited in the cross section of countries in Figure 2 and in the U.S. times series in Figure 7. Using market clearing conditions and the demand functions in equation (6), one can derive the relationship between relative labor shares and prices in this model is,

$$\frac{L_a}{L_n} = \left\{ \frac{w}{p_a(1+\nu)} + \frac{\bar{a}\nu p_a}{(1+\nu)A_a} \right\} \times \left\{ \frac{\nu w}{(1+\nu)A_n} - \frac{\nu \bar{a} p_a}{(1+\nu)A_n} \right\}^{-1}, \tag{14}$$

which is clearly not log-linear in relative labor shares and relative prices as the data suggests.

To further see the problems discussed above, we performed the following quantitative experiment. We solved this model over a range of A, A_a, and A_n generating differences in aggregate GDP per worker, agriculture GDP per worker, and non-agriculture GDP per worker similar to that in the data. The preference parameter ν is calibrated the same as in the text and the subsistence consumption parameter was calibrated to achieve a reasonable fit of labor and GDP shares in agriculture to the data.

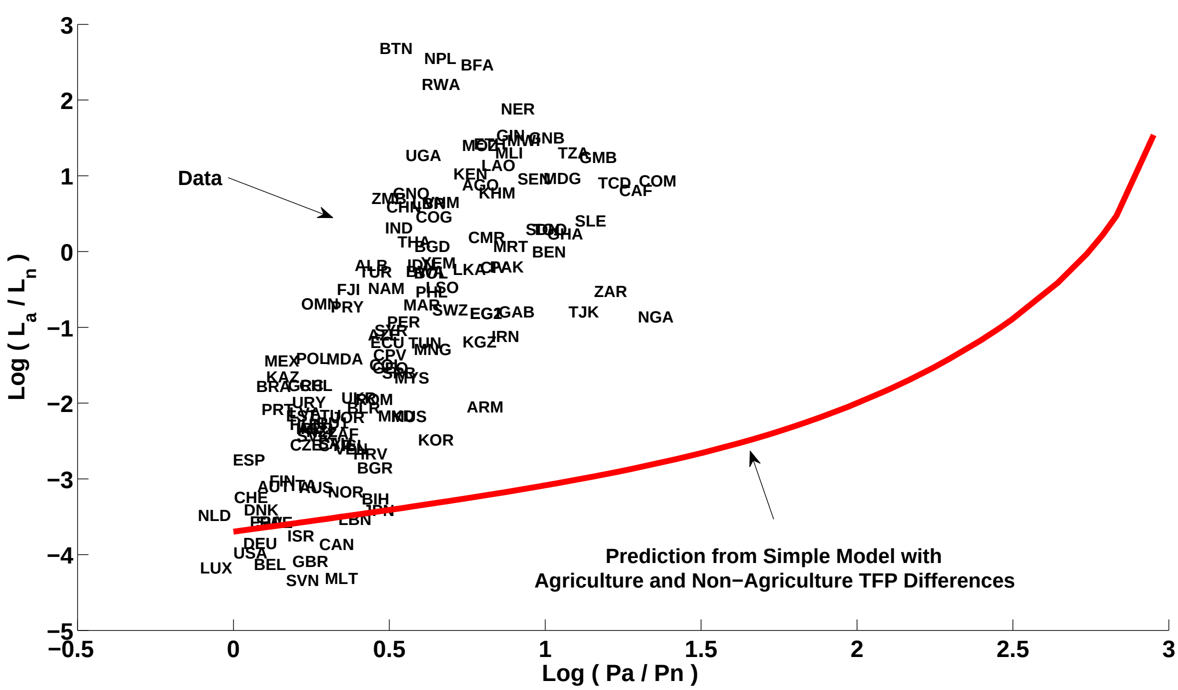

Figure 8: Relative Prices and Relative Labor in Data and Homogenous Labor Model

Figure 8 plots relative labor shares and relative price data and the predictions from the simple model with homogeneous labor. Confirming the discussion above, the simple model with homogenous labor over predicts the amount of variation in relative prices across countries. And the model does not generate a log-linear relationship between relative labor shares and relative prices as exhibited in the cross section of countries. We view these results as direct evidence against this particular model.

D Data Appendix

- **GDP Per Worker** — This data is from the Penn World Table version 6.2. series "rgdpch".
- **Labor Share in Agriculture** — This data comes from Table A.3 in the FAO Statistical Yearbook 2004 online edition.
- **Agriculture Share in GDP** — This data comes from Table G.1 in the FAO Statistical Yearbook online edition.
- **Relative Agriculture Prices** — This data is derived from author's calculations with original data from the World Bank's 2005 International Comparison Program online database. The sector "agriculture" is defined to be food and non-alcoholic beverages, alcoholic beverages and tobacco, codes (1101 and 1102). "Non-agriculture" is defined as all individual consumption, code (11), gross fixed investment, code (15), minus food, non-alcoholic beverages, alcoholic beverages and tobacco.
- **Relative Wages** — Wage data is from LABORSTA and the series wage data by economic activity is used. Agriculture corresponds directly with ISIC Revision 2 and 3 categories "Agriculture, hunting and forestry" and "Fishing".
- **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. 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.