

## Selection, Agriculture, and Cross-Country Productivity Differences<sup>†</sup>

By DAVID LAGAKOS AND MICHAEL E. WAUGH\*

*Cross-country labor productivity differences are larger in agriculture than in non-agriculture. We propose a new explanation for these patterns in which the self-selection of heterogeneous workers determines sector productivity. We formalize our theory in a general-equilibrium Roy model in which preferences feature a subsistence food requirement. In the model, subsistence requirements induce workers that are relatively unproductive at agricultural work to nonetheless select into the agriculture sector in poor countries. When parameterized, the model predicts that productivity differences are roughly twice as large in agriculture as non-agriculture even when countries differ by an economy-wide efficiency term that affects both sectors uniformly. (JEL J24, J31, J43, O11, O13, O40)*

Cross-country labor productivity differences are much larger in agriculture than in the non-agricultural sector (Caselli 2005; Restuccia, Yang, and Zhu 2008). Because developing countries have most of their workers in agriculture, their low productivity in agriculture accounts for nearly all of their low productivity in the aggregate. This implies that understanding why productivity differences in agriculture are so large compared to those of the non-agricultural sector is at the heart of understanding world income inequality.<sup>1</sup>

In this article we propose a new explanation for these productivity patterns in which the self-selection of heterogeneous workers determines sector productivity. We start from the well-known idea that in poor countries, where economy-wide efficiency is low, most people must work in the agricultural sector in order to satisfy subsistence consumption needs. This is what Schultz (1953) famously called the “food problem.” Our insight is that precisely because the majority of workers in

\*Lagakos: Arizona State University, 501 E. Orange Street, CPCOM 412A, Tempe, AZ 85287 (e-mail: [lagakos@asu.edu](mailto:lagakos@asu.edu)); Waugh: Department of Economics, Stern School of Business, New York University, 44 West Fourth Street, New York, NY 10012 (e-mail: [mwaugh@stern.nyu.edu](mailto:mwaugh@stern.nyu.edu)). For helpful comments we thank the three anonymous referees, Tasso Adamopoulos, David Atkin, Francesco Caselli, Berthold Herrendorf, Ayse Imrohoroğlu, Narayana Kocherlakota, Erzo Luttmer, Ellen McGrattan, Guillermo Ordoñez, Diego Restuccia, Andrés Rodríguez-Clare, Richard Rogerson, Jim Schmitz, Todd Schoellman, Jon Vogel, and seminar participants at ASU, BU, Chicago, 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 2010 ASU CASEE conference on Macroeconomics and Policy, the 2009 NBER Summer Institute Economic Growth Meeting, the 2009 SED Annual Meeting (Istanbul), the 2009 USC Macroeconomics Conference, and the 2009 SITE Growth and Development session. This version supersedes an earlier version with the title “Specialization, Agriculture, and Cross-Country Productivity Differences.”

<sup>†</sup> To view additional materials, visit the article page at <http://dx.doi.org/10.1257/aer.103.2.948>.

<sup>1</sup> Versions of this argument have been made by Caselli (2005); Restuccia, Yang, and Zhu (2008); Chanda and Dalgaard (2008); and Vollrath (2009).

poor countries are employed in agriculture, many of these workers must be relatively unproductive at agricultural work. In contrast, in rich countries, where economy-wide efficiency is high, those few workers selecting into agriculture must be those who are relatively most productive at agricultural work. Thus, two countries that differ in economy-wide efficiency will have even larger measured differences in agricultural productivity. By the same mechanism, they will have even smaller measured differences in non-agricultural productivity.

Our theory has two main ingredients. The first is nonhomothetic preferences, and in particular a subsistence consumption requirement in the agricultural good. This leads to an income elasticity of demand for agricultural goods less than one. The second ingredient is heterogeneity in individual (worker) productivity in each sector, combined with the assumption that workers choose where to supply their labor. This is the Roy (1951) model of self-selection based on comparative advantage. We combine these features into a two-sector general equilibrium version of the Roy model. Countries differ only in an economy-wide efficiency term; preferences and the distribution of individual productivity are taken to be identical across countries.

Within this economic environment, we provide a general condition on the heterogeneity in individual productivity that leads to productivity differences that are larger in agriculture than non-agriculture when countries differ only by an economy-wide efficiency term. The key condition is simple and economically meaningful: that comparative advantage aligns with absolute advantage. As long as workers who have a comparative advantage in a given sector have an absolute advantage (on average) in that sector, then our model qualitatively replicates the larger cross-country productivity differences in agriculture and smaller differences in non-agriculture.

To measure the quantitative importance of selection in explaining the sector productivity patterns at hand, we make flexible parametric assumptions on the distribution of individual productivity. In particular, we assume that sector productivities are drawn from dependent Fréchet distributions, where the dependence is captured parsimoniously using a copula. These assumptions allow us to calibrate the distribution parameters using moments computed from micro-level US wage data, specifically the variance of the nontransitory component of log wages in agriculture and non-agriculture, and the ratio of average wages in the two sectors.

Our benchmark quantitative experiment varies economy-wide efficiency of production to match the differences (of a factor 22) in aggregate GDP per worker between the ninetieth and tenth percentile countries of the world income distribution. It then computes the model's implications for differences in sector labor productivity, which arise only from selection. We find that the model predicts a factor 29 difference in agricultural productivity and a factor 13 difference in non-agricultural productivity, compared to factors of 45 and 4 in the data. In other words, the model produces roughly twice as much variation in agricultural productivity as non-agricultural productivity, compared to roughly ten times as much variation in the data. We conclude from this experiment that, while selection does not explain everything, it can lead to substantial amplification of large exogenous differences in economy-wide efficiency.

We then perform two alternative experiments that provide additional insight about the mechanics of the model's selection mechanism. The first alternative experiment varies economy-wide efficiency to match the differences (of a factor four) in

non-agricultural labor productivity between the tenth and ninetieth percentile countries. The model predicts that agricultural and aggregate productivity differences are only slightly larger than those of non-agriculture. The reason is that employment shares in agriculture are counterfactually similar between the rich and poor countries (3 percent versus 12 percent). This means that differences in the composition of workers by sector induced by selection are minimal and, hence, so are differences in sector productivity. These findings imply that selection forces by themselves are unable to amplify small exogenous differences in economy-wide efficiency into the much larger sector productivity differences observed in the data.

The second alternative experiment varies economy-wide efficiency to match the non-agricultural productivity differences across countries and varies an agriculture-specific efficiency term to match the employment share in agriculture in the tenth percentile country. The model predicts roughly ten times as much variation in agricultural productivity as non-agricultural productivity, roughly as in the data. Selection accounts for a factor of three of this difference. This experiment highlights the feature that selection forces can amplify efficiency differences of a sector-specific as well as a general nature, and that the quantitative importance of selection is substantial so long as the model matches the employment shares in agriculture in rich and poor countries.

We next extend the model to include capital and land and find that these forces increase the overall explanatory power of the model while leaving the importance of the selection channel largely unchanged (relative to the benchmark). When calibrated, the extended model produces four times as much variation in productivity in agriculture as non-agriculture. The improved performance comes from the well known role that land plays as a fixed factor in agriculture (e.g., the models of Restuccia, Yang, and Zhu 2008; Adamopoulos and Restuccia 2011; and Herrendorf and Teixeira 2011). While the importance of selection is similar in magnitude as in the benchmark model, decomposing the results into the contribution from land versus selection shows that selection is as important or more than the effects from land alone.

We find that in both the benchmark and the extended versions of the model, the quantitative predictions are consistent with other important features of the data not targeted directly. In particular, both predict a large wage gap between agricultural and non-agricultural workers, as in the data. This is in contrast to other papers in the literature, which reconcile this wage gap using some sort of exogenous barrier to workers moving out of agriculture (e.g., the work of Caselli and Coleman 2001; Restuccia, Yang, and Zhu 2008; Adamopoulos and Restuccia 2011; Tombe 2011; and Herrendorf and Teixeira 2011). Both are also quantitatively consistent with the higher employment shares in agriculture in poor countries, and the higher relative prices of agricultural goods in poor countries.

To illustrate how our theory works in practice, we provide one concrete example of how agricultural workers in developing countries are on average less productive at agricultural work than their counterparts in rich countries. Specifically, we cite evidence that women are less productive than men on average in agricultural work, and use cross-country data to document that women form a larger fraction of all agricultural workers in developing countries than in richer countries. Putting these together implies that poor countries have lower measured productivity in agriculture

in part because they employ more workers with relatively low productivity at agricultural work, just as our theory predicts.

Our article is the first to propose and assess the role of selection in understanding why productivity differences in agriculture are so much larger than in other sectors. This mechanism is distinct from previous explanations in the literature, most of which focus on frictions that are specific to the agricultural sector. For example, Restuccia, Yang, and Zhu (2008) argue that the larger productivity differences in agriculture are due partly to barriers to the adoption of intermediate goods in agriculture, Adamopoulos and Restuccia (2011) focus on the role of policies that misallocate farm land in developing countries, and Donovan (2011) argues that agricultural risk greatly reduces the incentives for farmers to use intermediates such as fertilizers.

One key difference is that our paper can reconcile some of the observed sector productivity patterns from distortions in poor countries which do not disproportionately affect agriculture, such as weak institutions, as emphasized by, e.g., Hall and Jones (1999) and Acemoglu, Johnson, and Robinson (2001, 2002). This suggests that policies that improve productivity in a general sense may disproportionately raise measured productivity in agriculture. Still, the selection mechanism in the current article is best thought of as complementary to other theories in the literature, in that it amplifies underlying productivity differences—either of a general or agriculture-specific nature—into even larger differences in measured agricultural productivity. This observation is important because it is unlikely that one story alone can completely explain why there is so much more productivity variation in agriculture than in other sectors, given the enormous magnitude of the difference.

## I. Motivating Evidence

In this section, we review the evidence that cross-country labor productivity differences are much larger in agriculture than in the non-agricultural sector. We then provide new calculations, and discuss existing evidence, suggesting that these sector labor productivity differences largely reflect sector differences in Total Factor Productivity (TFP).

Table 1 reproduces the findings of Caselli (2005), who constructs Purchasing Power Parity (PPP)-adjusted measures of labor productivity in the agricultural and non-agricultural sectors of 79 countries. His calculations combine PPP-adjusted GDP per worker data from the Penn World Tables with PPP-adjusted agricultural value-added-per-worker data from the Food and Agriculture Organization (FAO) constructed by Prasada Rao (1993).

The first row of Table 1 reports that the difference in aggregate output per worker between the ninetieth to tenth percentile of the world income distribution is a factor 22. In agriculture, this difference is a factor 45, while in non-agriculture it is a factor of just four. The last column shows that the ratio of agriculture to non-agriculture productivity differences is 10.7. In other words, there is more than ten times as much variation in agricultural productivity across countries than there is in non-agricultural productivity.<sup>2</sup>

<sup>2</sup>In independent work, Restuccia, Yang, and Zhu (2008) arrive at a very similar conclusion.

TABLE 1—SECTOR LABOR PRODUCTIVITY DIFFERENCES AND EMPLOYMENT SHARES

	Agriculture	Aggregate	Non-agriculture	Ag/non-agriculture ratio
90–10 labor productivity difference	45	22	4	10.7
Employment shares ninetieth percentile country	3		97	
Employment shares tenth percentile country	78		22	

*Notes:* The aggregate productivity difference is the ratio of GDP per worker between the ninetieth and tenth percentile countries. Sector productivity differences are the ratio of sector output per worker in the ninetieth and tenth percentile countries. The Ag/Non-agriculture Ratio is the agriculture productivity difference divided by the non-agriculture productivity difference.

*Source:* Caselli (2005).

The second and third rows report the percent of employment in agriculture in the ninetieth and tenth percentile countries. In the ninetieth percentile country, just 3 percent of labor is in agriculture, while the other 97 percent is in the non-agricultural sector. In the tenth percentile country, in contrast, 78 percent of workers are in agriculture, compared to 22 percent in non-agriculture. In short, a key distinction between rich and poor countries is that agriculture employs most people in the poorest countries and virtually nobody in the richest countries.

Simple accounting exercises show that the divide between agriculture and non-agriculture accounts for much of aggregate productivity differences. Caselli (2005) computes the hypothetical 90–10 ratio of aggregate output per worker by giving the agricultural productivity level of the ninetieth percentile country to all countries. He finds that the 90–10 ratio would be a factor of 1.6, down from the actual factor of 22! Similarly, by hypothetically giving an agricultural employment share of 3 percent, as in the ninetieth percentile country, to all countries, the 90–10 ratio would be just a factor 4.2.

One simple explanation of these sector labor productivity patterns is that developing countries use much less capital per worker in agriculture than in rich countries and use only modestly less capital per worker in non-agriculture. The main challenge to testing this hypothesis is the limited data on capital stocks by sector across countries. Caselli (2005) addresses this limitation by making the plausible assumption that rates of return to capital are equated across sectors and then using aggregate capital stock data to allocate capital to each sector. For a set of 65 countries for which comparisons can be made, he finds that capital explains 15 percent of cross-country productivity differences in agriculture, and 59 percent in non-agriculture. Thus, his calculations suggest capital differences are indeed important in both sectors, but there are still bigger residual productivity differences in agriculture even after taking capital into consideration.

To complement these findings, we conducted our own accounting exercises for a smaller set of countries using data on agricultural capital stocks constructed by Butzer, Mundlak, and Larson (2010). These data contain the values of machinery, equipment, livestock, and tree stock used in agricultural production in a set

of 28 countries from all income levels. As we detail in the online Appendix, we combine these data with estimates of the aggregate capital stocks constructed from the Penn World Tables to create estimates of the non-agricultural capital stocks in each country. The resulting sector capital data allow us to conduct accounting exercises in the same manner as Caselli (2005).

We find that using these new data, capital accounts for 22 percent of cross-country productivity differences in agriculture, and 29 percent in non-agriculture. Thus, these exercises largely corroborate the findings of Caselli (2005). While both sets of calculations have their limitations, both suggest that capital-per-worker differences are important in both sectors but unlikely to be the main cause of the larger differences in agricultural labor productivity across countries. In this sense, our findings are consistent with those of Chanda and Dalgaard (2008) and Vollrath (2009), who conclude that low agricultural (and aggregate) labor productivity in the developing world largely reflects their low TFP in agriculture.

## II. Model of Agricultural and Non-agricultural Productivity

In this section we formalize our model economy and characterize its equilibrium. The model predicts, under conditions that we describe, that exogenous differences in economy-wide efficiency across countries lead to even larger differences in agricultural productivity, and smaller differences in non-agricultural productivity.

### A. Preferences and Endowments

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

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

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 consumption requirement, and  $\nu$  governs the relative taste for non-agricultural consumption. These “Stone-Geary” preferences ensure that Engel’s Law holds, namely that the income elasticity of demand for food is less than one.

Each worker is endowed with one unit of time which she supplies inelastically to the labor market. Each worker is also endowed with a vector of “individual productivities,” denoted  $\{z_a^i, z_n^i\}$ , 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. The budget constraint of worker  $i$  is

$$(2) \quad p_a c_a^i + c_n^i \leq y^i,$$

where  $y^i$  is labor income (described in more detail below),  $p_a$  is the relative price of agriculture, and the non-agricultural good is taken to be the numeraire.



### B. Production

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

$$Y_a = A \cdot L_a \quad \text{and} \quad Y_n = A \cdot L_n,$$

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. Economies differ only in  $A$ , and we assume that each economy is closed.<sup>3</sup>

Let  $\Omega^a$  and  $\Omega^n$  denote the sets of workers 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

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

### C. Optimization and Equilibrium

An equilibrium of the economy consists of a relative agricultural price,  $p_a$ , wages per efficiency unit of labor in each sector,  $w_a$  and  $w_n$ , and allocations for all workers, such that all workers optimize and both labor markets 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.

Workers take prices and wages as given when they optimize. The problem for a worker is first to choose which sector to supply her labor, and then to maximize her utility, (i), subject to her budget constraint, (ii). Because of competition, the wages per efficiency unit of labor are

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

<sup>3</sup>The online Appendix discusses the implications of international trade for our results. See also Gollin, Lagakos, and Waugh (2011) for more on the open-economy implications of selection in multisector models, or Tombe (2011) and Teignier (2012) for more on how the lack of agriculture imports by countries with unproductive agriculture sectors can impede growth and structural change.

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

$$(3) \quad \frac{z_n^i}{z_a^i} \geq p_a.$$

Thus, the workers that enter non-agriculture are those whose productivity there 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 worker’s problem is standard, and optimal demands are

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

Due to the subsistence consumption requirement, workers consume relatively more agricultural goods when their income is lower. The lower is  $\nu$ , the higher is the ratio of agricultural to non-agricultural consumption.

#### D. Qualitative Features of Equilibrium

We now show that, in equilibrium, economy-wide efficiency determines the relative price of agriculture, which in turn determines the selection of workers and productivity in each sector.

*Relative Price of Agriculture Higher in Poorer Countries.*—The first important result is that, in equilibrium, the relative price of agriculture is higher in countries with lower economy-wide efficiency. We formalize this result as:

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

The proof is in the online Appendix. The intuition for this result is that a higher price of agricultural goods is needed in the poor economy in order to induce workers to work in the agricultural sector. To see this, let  $p_a^R$  be the equilibrium relative price in the rich economy. If  $p_a^R$  were the equilibrium price in the poor economy as well, then by (3), the sector labor-supply cutoffs would be the same in both countries, and so would the share of workers in agriculture. But because of the subsistence consumption requirement, the poor economy demands a relatively larger fraction of agricultural goods, and 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.

*Individual Productivity Distribution and Sectoral Productivity Differences.*—We now turn to the link between the distribution of individual productivity and sector aggregate productivity in equilibrium. Proposition 2 describes conditions on the



individual productivity distribution that are sufficient for economy-wide efficiency differences to lead to larger differences in agricultural labor productivity and smaller differences in non-agricultural labor productivity. See the online Appendix for the proof.

**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*

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

Intuitively, Proposition 2 says that as long as workers who have a comparative advantage in a given sector have an absolute advantage (on average) in that sector, then productivity differences will be larger in agriculture than non-agriculture across the two economies. The reason is as follows. As  $A$  rises, the relative price of agriculture falls (by Proposition 1), and only workers with a greater comparative advantage in agriculture (i.e., a higher  $z_a/z_n$  ratio) choose to work in agriculture. Then, since workers with a greater comparative advantage also have a greater absolute advantage, it follows that agricultural sector productivity increases. The second part of Proposition 2 says that, for a similar reason, non-agricultural productivity differences are smaller than  $A$  differences if workers with a greater comparative advantage in non-agriculture have a higher expected productivity in that sector.

Note that both heterogeneity in worker productivity and nonhomothetic preferences are necessary for Proposition 2 to hold. When all workers are identical in productivity, then changes in  $A$  induce changes in the share of workers in agriculture but (trivially) do not change the average individual productivity by sector. When preferences are homothetic, relative prices and, hence, the allocation of workers by sector are independent of  $A$ . Thus, in each sector, average individual productivity is identical across countries.

In general, at least one of the conditions of Proposition 2 must hold (see Heckman and Honore 1990). Thus, at the very least, our theory qualitatively delivers productivity differences in one sector that differ from the aggregate in a way consistent with the data. Of course, it can also explain the patterns of both sectors. 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.<sup>4</sup>

### E. Independent Fréchet Individual Productivities

In this section, we illustrate the mechanics of the theory using a version of the model which assumes independent Fréchet distributions on individual productivity.

<sup>4</sup>One can show that both conditions of Proposition 2 hold whenever individual productivities 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 of independent Fréchet distributions that we focus on below.

This example helps demonstrate how the size of the mechanism’s effects depend on (i) the variance in individual productivity, and (ii) differences in the sector employment shares across the economies being compared. Furthermore, this example is a special case of the individual productivity distribution used for quantitative analysis in Section III.

ASSUMPTION 1: Let  $z_a$  and  $z_n$  be drawn independently from Fréchet distributions:

$$G(z_a) = e^{-z_a^{-\theta}} \quad \text{and} \quad G(z_n) = e^{-z_n^{-\theta}}, \quad \text{with } \theta > 1.$$

The parameter  $\theta$  controls the dispersion of individual productivity in each sector, with a smaller  $\theta$  implying more productivity dispersion across individuals and a higher  $\theta$  meaning less dispersion.<sup>5</sup> This distributional assumption conveniently relates equilibrium employment shares in agriculture, the relative price of agriculture, and parameter  $\theta$ . The equilibrium share of workers in agriculture is

$$(4) \quad \pi_a = \text{Prob} \{Az_n^i \leq p_a Az_a^i\} = \frac{1}{p_a^{-\theta} + 1}.$$

By (4), one can see that as  $p_a$  rises, the share of workers in agriculture rises as well. Furthermore, the responsiveness of the share of workers in agriculture to  $p_a$  depends on the productivity-dispersion parameter  $\theta$ . Manipulating (4), and a similar equation for non-agriculture, yields a log-linear relationship in the ratio of the agricultural to non-agricultural worker shares ( $\pi_n$ ) and the relative price of agricultural goods:

$$(5) \quad \log(\pi_a/\pi_n) = \theta \log(p_a).$$

Intuitively, with a low  $\theta$ , 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  $\theta$ , meaning small productivity dispersion, implies that only small changes in the relative price of agriculture are needed to induce workers to switch sectors.

Both conditions on the productivity distribution in 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. To see this note that expected individual productivity in the two sectors are

$$(6) \quad E(z_a|z_a/z_n > 1/p_a) = \gamma \pi_a^{-1/\theta}, \quad \text{and} \quad E(z_n|z_n/z_a > p_a) = \gamma \pi_n^{-1/\theta},$$

where the constant  $\gamma$  is the Gamma function evaluated at  $(\theta - 1)/\theta$ . Equation (6) relates expected individual productivity to the share of workers in a sector and through equation (5) the relative price. A decrease in the relative price of agriculture decreases the share of workers in agriculture. This then leaves a more selected set

<sup>5</sup>This distribution has been used by Eaton and Kortum (2002) and others to analytically solve multicountry 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.

of workers in agriculture with higher average agricultural productivity. Similarly, because the share of workers increases in non-agriculture, non-agricultural productivity decreases. The magnitude of these changes depends on the parameter  $\theta$ .

Differences in  $A$  across economies will lead to relative price differences (Proposition 1). This then leads to differences in employment shares (equation (5)) and, hence, to larger productivity differences in agriculture and smaller ones in non-agriculture across economies. We formalize this as:

**COROLLARY 1:** *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*

$$(7) \quad \frac{Y_a^R/N_a^R}{Y_a^P/N_a^P} = \left( \frac{\pi_a^P}{\pi_a^R} \right)^{\frac{1}{\theta}} \left( \frac{A^R}{A^P} \right) > \frac{A^R}{A^P},$$

$$\text{and} \quad \frac{Y_n^R/N_n^R}{Y_n^P/N_n^P} = \left( \frac{\pi_n^P}{\pi_n^R} \right)^{\frac{1}{\theta}} \left( \frac{A^R}{A^P} \right) < \frac{A^R}{A^P}.$$

See the online Appendix for the proof. Dispersion in individual productivity controls the magnitude of the sector productivity difference from the aggregate. A lower  $\theta$  leads agricultural productivity differences to be larger than the aggregate (since  $\pi_a^R < \pi_a^P$  in equilibrium). As  $\theta$  approaches infinity, heterogeneity in individual productivity disappears, selection effects are eliminated, and the ratio of agriculture productivity converges downward toward the aggregate productivity ratio. A similar argument illustrates that the non-agricultural productivity difference is smaller than the difference in  $A$ , with the magnitude of the difference again shrinking to zero as individual productivity dispersion is reduced to zero.

With large cross-country differences in employment shares these effects can be potent. Consider, for example, the countries at the tenth and ninetieth percentile of the world income distribution, which have 78 and 3 percent of their work force in agriculture. With a  $\theta$  of five, these differences in agricultural employment shares amplify the underlying  $A$  differences by a factor of two. In contrast, for countries that do not differ dramatically in agricultural employment shares, these effects will be more modest.

### III. Quantitative Analysis

We now present a richer model that we calibrate and use to assess the quantitative importance of the theory. This model differs from the version of the previous section by allowing for correlation between individual productivity draws and different degrees of productivity dispersion in the two sectors. Introducing these richer features is important because it allows for greater flexibility in matching the data, and because it allows our theory to fail, in the sense that nothing assures that both assumptions on the individual productivity distribution in Proposition 2 hold.

We then calibrate the parameters of the model to US data and then lower economy-wide efficiency to study the implications for productivity by sector in developing countries. Our calibration strategy is to use moments from the distribution of wages in the United States, namely the nontransitory component of the variance of log wages in each sector and the ratio of sector average wages, to discipline the distribution parameters. Employment statistics in agriculture in the United States discipline preference parameters. Aggregate productivity differences across countries discipline the extent to which we lower economy-wide efficiency in our main experiment.

### A. Dependent Fréchet Individual Productivity Distribution

We set the joint distribution of individual productivities to be

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

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

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

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  $H(z_n)$ .<sup>6</sup> The parameter  $\rho \in (-\infty, \infty) \setminus \{0\}$  determines the extent of dependence, with a positive (negative) value of  $\rho$  representing positive (negative) dependence between the draws.<sup>7</sup> The marginal distributions themselves are Fréchet, with dispersion parameters  $\theta_a$  and  $\theta_n$ , and scale parameters normalized to one. The lower are  $\theta_a$  and  $\theta_n$ , the higher is the variation in individual productivity in agriculture and non-agriculture.

We choose this distribution for three main reasons. First, it allows for richness and flexibility in matching the data, while maintaining a parsimonious set of parameters. One dimension of richness relative to the simple version of Section IIE, is that individual productivity draws are no longer independent across sectors. This allows for the feature that individuals who are relatively productive in one sector are relatively productive in the other sector as well. The other dimension of richness is that dispersion in individual productivity is no longer the same in each sector. Since non-agricultural 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 distribution allows for this possibility.

The second reason for this choice of distribution is that Fréchet marginals in each sector contain a sensible economic interpretation, which is as follows. The Fréchet distribution is an extreme-value distribution, representing the distribution of the

<sup>6</sup>A copula is a function that allows for the creation of multivariate distributions out of arbitrary univariate distributions; e.g., Nelsen (2006). The Frank copula generates dependence between draws that is *radially symmetric*, i.e., not systematically stronger when closer to the right or left tails of the distribution. Other copulas, such as the Clayton or the Gumbel copula, do not have this feature.

<sup>7</sup>When  $\rho = 0$ ,  $C[u, v]$  is defined as  $u \cdot v$ .

maximum of independent draws from some underlying distribution.<sup>8</sup> Thus, the draw  $z_n^i$  can be thought of as the maximum of worker  $i$ 's individual productivity draws in a large set of distinct non-agricultural tasks. A similar interpretation can be given to  $z_a^i$ .<sup>9</sup>

The third reason for choosing this distribution is that it allows our theory to fail. In particular, there is nothing inherent in this distribution that assures that both assumptions on the individual productivity distribution in Proposition 2 hold. Whether both conditions hold will be dictated by the data and the calibration procedure.

### B. Calibration of Individual Productivities

To calibrate the individual-productivity distribution parameters, our strategy uses micro-level wage data from the United States. Formally, we jointly calibrate  $\theta_a$ ,  $\theta_n$ , and  $\rho$  to match three moments: the variance of the nontransitory component of log wages in agriculture and non-agriculture (which we define below) and the ratio of average wages in the two sectors.

While all three parameters are jointly determined, each has an intuitive relationship with one of the moments picked. The parameters  $\theta_a$  and  $\theta_n$  are disciplined by variation in the nontransitory component of log wages in the two sectors. Because a worker's wage in the model equals the value of her marginal product, variation in individual productivity maps into variation in wages across workers.

The dependence parameter  $\rho$  is disciplined by the ratio of average wages in agriculture to average wages in non-agriculture, with a lower ratio implying a higher  $\rho$ . The intuition is as follows. For high values of  $\rho$ , workers tend to get either two high draws or two low ones. Because of the higher variance in non-agricultural productivity (implied by the calibration procedure, as we explain below), those with the high draws are more likely to have a comparative advantage in non-agriculture. This implies that most of the high-wage workers are in the non-agricultural sector, and that the ratio of average agricultural wages to average non-agricultural wages is low. For low values of  $\rho$ , in contrast, each sector employs workers with high sector-specific skills, and higher wage individuals are more equally distributed across sectors. Hence, the ratio of average agricultural wages to non-agricultural wages is higher.

Our wage data come from the US Current Population Survey (CPS) for 1996 through 2010.<sup>10</sup> Our sample includes all individuals who have nonmissing data on income and hours worked, including both self-employed and salaried workers, whose wage is at least the federal minimum wage, and who can be matched across years (as we describe below). We calculate each individual's wage as her labor income in the previous year divided by her hours worked in the previous year. We define labor income as the sum of salary income plus  $0.66 \times$  business income plus  $0.46 \times$  farm

<sup>8</sup>By the extreme value theorem, the maximum of independent draws from any distribution converges in distribution (once properly normalized) to one of three extreme value distributions: the Fréchet, the Gumbel, or the Weibull.

<sup>9</sup>Yet another advantage of Fréchet distributions is that they produce wage distributions with fat right tails, as in the data, while other prominent distributions fail in this dimension. For example, we find that a version of our model with log normal individual-productivity distributions generates tails that are too thin compared to the data. Heckman and Sedlacek (1985) arrive at a similar conclusion. Details of our calculations are available on request.

<sup>10</sup>The online Appendix provides a more detailed description of the data and strategy for estimating the variance of the nontransitory component of wages by sector.

income, where the fractions of business and farm income assigned to labor income are taken from Valentinyi and Herrendorf (2008). We define agricultural workers to be those whose primary industry of employment is agriculture, forestry, or fishing, and non-agricultural workers to be all other workers.

We measure the variance of the nontransitory component of wages by sector as follows. The short panel dimension of the CPS allows us to match a subset of individuals across consecutive years. For these individuals, we estimate the variance of the nontransitory component of wages in each sector as the covariance of the two wage observations across individuals. The online Appendix provides a detailed description of the data and our estimation approach. We end up with estimates of the variance of the nontransitory component of log wages of 0.144 in agriculture and 0.224 in non-agriculture, which we target in our calibration.

The final moment we target in our calibration is the ratio of average wages in agriculture to average wages in non-agriculture. Using the CPS data, we calculate this ratio to be 0.701.<sup>11</sup>

These moments imply parameter values of  $\theta_a = 5.3$ ,  $\theta_n = 2.7$ , and  $\rho = 3.5$ . The estimates of  $\theta_a$  and  $\theta_n$  mean that there is more variation in individual productivity in non-agricultural work than in agricultural work, which seems reasonable given that non-agricultural work encompasses more types of economic activities. While  $\rho$  itself is hard to interpret, the associated Spearman rank correlation coefficient is 0.35 and the linear correlation coefficient is 0.44. Thus, the calibrated model features a modest positive correlation in individual productivities.

### C. Calibration of Preference Parameters

Given the calibrated values of  $\theta_a$ ,  $\theta_n$ , and  $\rho$ , we pick  $\bar{a}$  and  $\nu$  to jointly match two moments from US data. The first moment we target is the fraction of workers in agriculture from US data, which is just below 2 percent. The second is a long-run agricultural employment share of 0.5 percent, which has been used by others in the literature, in particular Restuccia, Yang, and Zhu (2008). The resulting values for the minimum consumption requirement and non-agriculture preference parameters are  $\bar{a} = 2.43$  and  $\nu = 276$ .

The resulting parameter  $\bar{a}$  is consistent with independent estimates of the size of the subsistence consumption requirement in developing countries. Rosenzweig and Wolpin (1993) and Atkeson and Ogaki (1996), both of which use panel data from a sample of rural households in India, estimate a subsistence consumption need of around 33 percent of the average income of Indian villagers. When we compute the subsistence consumption requirement in our model economy with  $A$  calibrated to mimic India's per capita GDP relative to the United States we find that  $\bar{a}$  is 27 percent of average income.

<sup>11</sup> Herrendorf and Schoellman (2011) find that the low average wage of agriculture workers in the United States is accounted for largely by lower measured returns to both schooling and experience among agriculture workers than non-agriculture workers. One interpretation of this finding is a selection story such as the one presented here.



TABLE 2—90–10 PRODUCTIVITY DIFFERENCES, DATA AND BENCHMARK MODEL

	Agriculture	Aggregate	Non-agriculture	Ag/non-agriculture ratio
Data	45	22	4	10.7
Model	29	22	13	2.2
Without selection	19	19	19	1.0

*Notes:* The aggregate productivity difference is the ratio of GDP per worker between the ninetieth and tenth percentile countries. Sector productivity differences are the ratio of sector output per worker in the ninetieth and tenth percentile countries. The ag/non-agriculture ratios are the agriculture productivity differences divided by the non-agriculture productivity differences.

*Source:* Authors' calculations and Caselli (2005).

#### D. Quantitative Predictions for Sector Productivity Differences

To explore the quantitative implications of our model, we perform the following experiment. Beginning with a value of  $A$  normalized to one for the benchmark economy (calibrated to the United States), we lower  $A$  to match GDP per worker for a country in the ninetieth percentile of the income distribution, and then for a country in the tenth percentile. We then compare the model's predictions for sector labor productivity differences between the tenth and ninetieth percentile countries to those of the data.

Table 2 shows the results of the experiment. By construction, the gap in aggregate labor productivity is a factor 22 in both the model and data. This gap is generated with a difference in  $A$  of a factor 19.<sup>12</sup> At the sector level, the model predicts a factor 29 difference in agricultural productivity, and a factor 13 difference in non-agriculture. In the data, the differences are a factor 45 in agriculture and a factor 4 in non-agriculture. Thus, the selection channel in the model generates quantitatively large differences between sector and aggregate productivity differences, but not quite as large as in the data.

To provide a more concrete metric for the overall quantitative importance of selection, the last column of Table 2 shows the ratio of the productivity differences in agriculture to those of non-agriculture in the model, which is 2.2. The implication is that if selection were the only phenomenon at work, agricultural productivity differences would be roughly twice as large as productivity differences in non-agriculture. The equivalent figure in the data is 10.7. For illustration, the bottom row presents the model's predictions without selection, i.e., when worker heterogeneity is shut down. In this case the ratio is 1.0, as the productivity gaps are the same in each sector as the  $A$  differences themselves.

Table 3 provides more insight about where the selection effects come from. For each country, the table reports the expected individual productivity of workers in each sector relative to the population mean (unconditional expected productivity). In the ninetieth percentile country, the average agricultural worker has productivity in agriculture 1.55 times the population mean. Recall that the ninetieth percentile country in the model has just 3 percent of workers in agriculture. What the model

<sup>12</sup> Aggregate labor productivity is expressed as GDP per worker at prices of the rich country; we find similar results when using Gheary-Khamis international prices.

TABLE 3—90–10 EXPECTED INDIVIDUAL PRODUCTIVITY RELATIVE TO POPULATION MEAN

Country	Agriculture	Non-agriculture
Ninetieth percentile	1.55	1.01
Tenth percentile	1.00	1.42
90/10 ratio	1.55	0.71

*Note:* Expected individual productivity relative to the population means are calculated as  $E(z_a | z_a/z_n > 1/p_a)/E(z_a)$  and  $E(z_n | z_n/z_a > p_a)/E(z_n)$  and represent the conditional average individual productivity in agriculture and non-agriculture relative to the unconditional average productivities.

*Source:* Authors' calculations.

predicts is that this small set of workers are in fact much more productive in agriculture than a worker taken at random from the population. In the tenth percentile country, in contrast, agricultural workers have productivity roughly equal to the population mean, which is natural given that the majority of workers are in agriculture. The ratio of average productivity of agricultural workers in the two countries is 1.55. Note that this corresponds to the ratio of the agricultural productivity difference to the underlying  $A$  difference in Table 2 (i.e., 29/19).

In non-agriculture, selection forces work in the opposite direction. In the ninetieth percentile country, non-agricultural workers have productivity just 1.01 times the population mean. This is not surprising, as virtually all workers are employed in the non-agricultural sector in this country. In the tenth percentile country, with roughly one third of workers employed in the non-agricultural sector, non-agricultural workers have productivity 1.42 times the population mean. Taking a ratio of the ninetieth to tenth percentile of the country income distribution gives 0.71, which corresponds to the ratio of the non-agricultural productivity difference to the underlying  $A$  difference in Table 2 (i.e., 13/19).

These observations imply that workers with a comparative advantage in agriculture (non-agriculture) also have an absolute advantage in agriculture (non-agriculture). This means that both conditions on the individual-productivity distribution of Proposition 2 hold. We note that there was nothing in our calibration strategy that guaranteed this outcome. Indeed, there exist parameter combinations for which one of the conditions fails. The sensitivity analysis of Section III G provides one such example and illustrates the dimensions on which it is counterfactual to the data.

### *E. Assessment of Calibrated Model's Cross-Country Implications*

The model has a variety of other predictions for the cross section of countries. Below we highlight some of its other main quantitative implications and compare them to cross-country data.

*Agricultural Wage Gaps.*—One novel prediction of our model is that average wages are much lower in agriculture than non-agriculture even though there are no barriers to workers moving between sectors (as in the models of Caselli and Coleman 2001; Restuccia, Yang, and Zhu 2008; Herrendorf and Teixeira 2011; Adamopoulos and Restuccia 2011; and Tombe 2011). In our model, this agricultural wage gap is

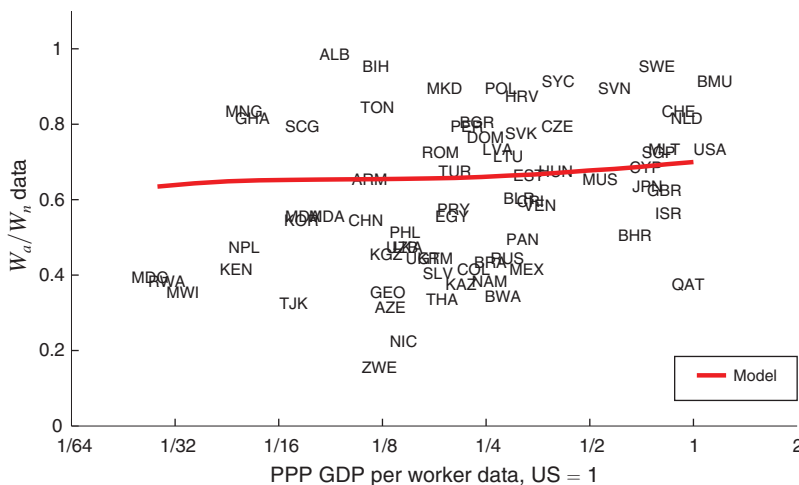


FIGURE 1. AVERAGE WAGE IN AGRICULTURE RELATIVE TO NON-AGRICULTURE, DATA, AND MODEL

driven entirely by selection: in equilibrium, most of the high-wage individuals are those who possess a comparative advantage in non-agricultural production and self-select into that sector (see Section IIIB).

Figure 1 plots the ratio of average wages in agriculture to non-agriculture against GDP per worker using wage data from the International Labor Organization (ILO) and the predictions from our model. In the data, virtually all countries have a ratio of agriculture wages to non-agriculture wages *below one*. The model also predicts a ratio below one. Furthermore, the model also predicts a modest increase in the wage ratios with GDP per worker, as in the data, albeit with a less steep increase than in the data.

The success of our model on this dimension is important because new evidence suggests several prominent explanations are unable to account for gaps in income or value added per worker by sector. Using new data from a large set of developing countries, Gollin, Lagakos, and Waugh (2012) find that even after adjusting for sector differences in hours worked per worker, human capital per worker, and cost-of-living differences between agricultural and other areas, there are still large residual gaps in income and value added per worker in agriculture in developing countries. The current article suggests one potential explanation for the residual gaps, which is that selection forces induce higher wage workers to be disproportionately in the non-agricultural sector.

*Share of Workers in Agriculture.*—The model’s nonhomothetic preferences assure that the share of workers in agriculture declines in income per capita. Here we ask whether the model’s quantitative implications on this dimension are consistent with the data, as is needed to accurately gauge the importance of the selection mechanism (e.g., (8) in the analytical version of the model). Figure 2 plots data on the percent of employment in agriculture against GDP per worker data and the predictions from our calibrated model. Both the model and data predict a strong decline in

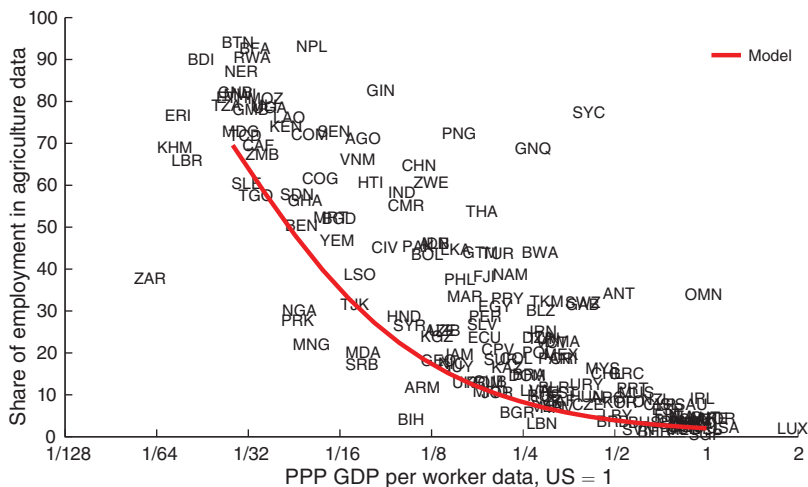


FIGURE 2. SHARE OF EMPLOYMENT IN AGRICULTURE, DATA, AND MODEL

agriculture’s share of employment in GDP per capita, albeit a somewhat stronger relationship in the data than in the model. In the data, the country in the tenth percentile of the income distribution has an employment share in agriculture of 78 percent, compared to 58 percent in the model. Both the data and model have a 3 percent share in agriculture in the ninetieth percentile country as per the calibration.

*The Relative Price of Agriculture.*—As Proposition 1 shows, the model predicts that relative agricultural prices are higher in poor countries than in rich countries. Figure 3 plots the predictions of our quantitative model, as well as data on the relative price of agriculture and GDP per worker. Our data on relative agricultural prices are constructed using 2005 data from the International Comparison Programme (ICP); the online Appendix provides the complete details. Figure 3 shows that relative agricultural prices systematically decline in GDP per worker, with a ratio of relative prices between countries in the ninetieth and tenth percentiles of GDP per worker of 2.5. The solid line in Figure 3 plots the model’s prediction. In the model, relative agricultural prices also systematically decline with GDP per worker, and the ratio between the ninetieth and tenth percentiles is 2.3.<sup>13</sup>

One concern is that the data are based on the prices that consumers pay for goods, not the price that producers receive. This distinction would reflect distribution margins that are not in the model. If distribution margins vary systematically with the level of development (see, for example, Adamopoulos 2009), then the relationship

<sup>13</sup>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 (2012), show that when partitioning International Comparison Program goods into agricultural and non-agricultural goods, the relative price of agriculture is higher in poor countries. They also show that partitioning goods into tradable and nontradable goods implies higher relative prices of tradables in poor countries, and partitioning goods into consumption and investment goods implies a higher relative price of investment goods in poor countries.

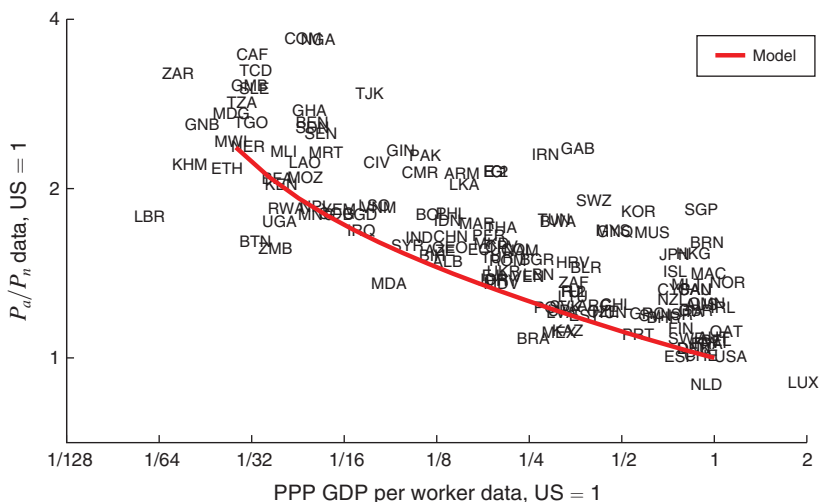


FIGURE 3. RELATIVE PRICE OF AGRICULTURAL GOODS, DATA, AND MODEL

in Figure 3 may not reflect differences in relative agricultural-producer prices. To address this concern, we examined relative agricultural-price data using producer prices constructed by Restuccia, Yang, and Zhu (2008). We find that, by these measures, relative agricultural prices systematically decline in GDP per worker—as our model predicts—and, in fact, the relationship is even stronger than for consumer prices. The online Appendix presents these findings in more detail.

F. Evidence Using Proxies for Individual Productivities

In this section we take a different approach to assessing the plausibility of the calibrated model. In particular, we use two plausible proxies for agricultural and non-agricultural individual productivity that are observable independent of the sector the worker selects, and provide evidence supporting key implications of our model.

The two proxies we use are height for agriculture and cognitive ability scores for non-agriculture. The rationale is that height reflects the “physical vigor” (Steckel 1995) useful in physically demanding jobs such as agricultural work (see Pitt, Rosenzweig, and Hassan 2010; Pitt, Rosenzweig, and Hassan 1990; Steckel 1995; and Strauss and Thomas 1998 plus the references therein). Cognitive ability scores in turn plausibly capture the verbal or other nonphysical capabilities often valued in non-agricultural activities (e.g., Case and Paxson 2008 or Miguel and Hamory 2009). While these proxies are certainly crude, they offer the advantage of being observable whether or not someone works in a particular sector, and have been (reasonably) widely measured in practice.

Using these proxies, we can compare the model’s correlation between individual productivities and the correlation between the observed proxies. As discussed in Section IIIB, the model’s correlation is positive, with a linear correlation coefficient of 0.44. This correlation is broadly consistent with the correlations between height and cognitive ability scores. Existing studies find correlations between height

and cognitive ability in the range of 0.10 to 0.30 (see Case and Paxson 2008 and the references therein). Thus, while the model's correlation overpredicts the data somewhat, it is consistent with the fact that the correlation in the data is positive yet modest in magnitude.

Another important implication of the model is that agricultural workers in rich countries are more productive in agriculture than the average worker (e.g., Table 3). Using height as a proxy for agricultural productivity, we should find that agricultural workers in rich countries are taller than the average worker. To check this prediction, we draw on height data for US adults collected in the 2009 National Health Interview Survey, conducted by the Centers for Disease Control (CDC). We find that the average agricultural worker is 172.4 cm tall, while the average worker is 170.0 cm tall. The difference of 2.4 cm, or roughly one inch, is statistically significant at well below the 1 percent confidence level. Furthermore, it is economically significant: according to the CDC, this difference is equivalent in magnitude to the overall increase in average height in the United States from 1960 to 2009.

The developing-country analog is that non-agricultural workers in poor countries are more productive in non-agricultural tasks than the average worker (again see Table 3). Using cognitive ability scores as a proxy for non-agricultural productivity, we should find that non-agricultural workers in poor countries have higher cognitive ability scores than average. While cognitive ability score data from developing countries are limited, the available evidence supports this implication. Using a unique dataset from Kenya, Miguel and Hamory (2009) find that among rural Kenyan students, cognitive ability scores are a very strong predictor of who later migrates out of agricultural areas to take non-agricultural employment. Their estimates suggest that students who score one standard deviation higher on cognitive ability tests are roughly 17 percent more likely to migrate out of agriculture areas after finishing school. In addition, other studies have found that those with greater schooling attainment are far more likely to migrate to non-agricultural areas (e.g., Lanzona 1998 and Beegle, De Weerd, and Dercon 2011). As schooling attainment is correlated with cognitive abilities, this evidence also supports the model's predictions that non-agricultural workers in developing countries have higher cognitive ability than average.

We conclude that, when using height and cognitive ability scores as proxies for agricultural and non-agricultural individual productivity, the available evidence is, in fact, consistent with the model's predictions. In particular, the correlation between the proxies does appear to be positive but modest, agricultural workers in the United States do appear to be selected on height, and non-agricultural workers in developing countries do appear to be selected on cognitive ability. Of course given the crudeness of these proxies and limited availability of data, we take this evidence as supportive rather than definitive.

#### *G. Sensitivity Analysis: Size of Correlation between Individual Productivities*

In this section we study the sensitivity of our results to the size of the correlation between individual productivity draws. This is an important issue because the correlation parameter helps determine the magnitude of the selection effects and when the conditions in Proposition 2 hold or not. Heckman and Honore (1990)



TABLE 4—SENSITIVITY OF SECTOR PRODUCTIVITY DIFFERENCES TO CORRELATION PARAMETER

	Correlation						
	0.00	0.20	0.30	0.35*	0.40	0.50	0.99
Ratio of average wage $\bar{w}_a/\bar{w}_n$	0.79	0.78	0.74	0.70*	0.66	0.61	0.52
Agriculture productivity difference	37	33	31	29*	28	26	21
Non-agriculture productivity difference	10	11	13	13*	14	15	18
Ag/non-agriculture ratio	3.8	3.0	2.5	2.2*	1.9	1.7	1.2

*Notes:* Results replicate the experiment of Table 2 but vary the correlation parameter. All other parameters are recalibrated as described in Section IIIC. The first row reports the Spearman rank correlation coefficient. Stars indicate the benchmark calibration. Sector productivity differences are the ratio of sector output per worker in the ninetieth and tenth percentile countries. The ag/non-agriculture ratios are the agriculture productivity ratios divided by the non-agriculture productivity ratios.

*Source:* Authors' calculations.

formalize this last point by showing that in the Roy model, the correlation in individual productivities determines how average productivity of workers in each sector relates to the unconditional averages, and in turn how comparative advantage aligns with absolute advantage, i.e., the conditions in Proposition 2.<sup>14</sup> Thus, we explore how varying the correlation affects our results.

To explore these issues, we recompute the results of our main experiment (of Table 2) under a range of correlation coefficients running from 0.00 (independence) to 0.99 (near perfect correlation) by varying the dependence parameter  $\rho$ . In each case, we recalibrate  $\theta_a$  and  $\theta_n$  to be consistent with the variance of the non-transitory component of log wages described in Section IIIB and recalibrate the preference parameters as described in Section IIIC. We do not attempt to match the ratio of average wages (since by varying  $\rho$ , we are no longer able to) but instead report the model's prediction for the sector wage ratio for each correlation value.

Table 4 shows the results of varying the model's correlation parameter in individual productivity, with the calibrated model in the center column (and marked with stars). The first row reports the Spearman rank correlation coefficient in each experiment. The second row reports the ratio of average wages in agriculture to non-agriculture. The third and fourth rows report the productivity differences between the ninetieth and tenth percentile countries in the two sectors. The final row presents the ratio of sector productivity differences.

One prominent feature in Table 4 is that higher values of correlation in individual productivity lead to smaller quantitative effects of selection. Starting with the calibrated model and increasing the correlation to 0.4 and 0.5 leads to agricultural gaps of 28 and 26, down from 29 in the calibrated model. Non-agricultural gaps rise to 14 and 15 up from 13. Thus, the model performs modestly worse in this range, with the combined effect of selection falling to a ratio 1.9 and 1.7 respectively. One challenge to correlation parameters in this range is the ratio of average sector wages is counterfactually low at 0.66 and 0.61, respectively.

<sup>14</sup> Heckman and Honore (1990) refer to the case when comparative advantage aligns with absolute advantage as the "standard case," and the case when agents with a comparative advantage have an absolute disadvantage as the "non-standard case."

In contrast, lower values of the correlation parameter lead to larger quantitative effects. Lowering the correlation to 0.3 and 0.2 leads to larger agricultural gaps of 31 and 33, and smaller non-agricultural gaps of 13 and 11. The combined effects rise to ratios of 2.5 and 3.0. The ratio of average wages also rises above the level found in the data, to 0.74 and 0.78.

The first and last data columns present some extreme examples of correlation, namely, no correlation and near-perfect (0.99) correlation. In the zero-correlation case, the model performs better than in the benchmark case, with agricultural and non-agricultural gaps of 37 and 10, and an overall ratio of 3.8. The wage ratio is counterfactually high at 0.79. In the case of near-perfect correlation, the agricultural gaps are a factor 21, while non-agricultural gaps are a factor 18, leading to an overall ratio of 1.2. Note that in this case, the underlying  $A$  differences are a factor of 22, which are larger than the agricultural productivity differences. The reason that the agriculture sector “flips” here is that in this case workers with a comparative advantage in agriculture have an absolute disadvantage there, i.e., one of the conditions in Proposition 2 does not hold. Thus, selection works in the opposite way as in the standard case. Of course a limitation of having such a high correlation is that the average wage in agriculture relative to non-agriculture is counterfactually low at 0.52.

#### IV. Alternative Quantitative Experiments

In this section we present two alternative quantitative experiments which provide additional insight about the mechanics of our model. The experiments demonstrate that the quantitative importance of selection is largest for economies with substantial differences in employment shares by sector, and that to get large employment share differences, the model requires large exogenous efficiency differences.

##### *A. Alternative Experiment No. 1: Calibrate to Non-agricultural Productivity Gaps*

The first alternative experiment lowers economy-wide efficiency differences to match non-agricultural productivity difference between the ninetieth and tenth percentile countries, which is a factor of four.

Table 5 illustrates the model’s predictions for sector labor productivity differences to those of the data. Note that the gap in non-agricultural productivity is a factor of four in the model (as in the data) by construction. The difference in aggregate productivity is a factor 4.3, substantially less than in the data, and agricultural productivity is now a factor of 5.5, far below the 45 in the data. Overall, the model predicts that agricultural and aggregate productivity differences are only slightly larger than those of non-agriculture, and selection forces are playing only a small role.

Why does the model fare so poorly in this case? The reason is that the small economy-wide efficiency differences required to match a non-agriculture productivity gap of four leave few workers in the poor country working in agriculture. This implies that employment shares in agriculture between the ninetieth and tenth countries in the model are virtually the same, with an 11 percent agricultural employment share in the tenth percentile country (rather than 78 percent in the data), compared to 3 percent in the ninetieth percentile country. Thus, agricultural workers are highly

TABLE 5—90–10 PRODUCTIVITY DIFFERENCES, ALTERNATIVE EXPERIMENT 1

	Agriculture	Aggregate	Non-agriculture	Ag/non-agriculture ratio
Data	45	22	4	10.7
Model	5.5	4.3	4	1.4
Without selection	4.3	4.3	4.3	1.0

*Notes:* This table displays results with  $A$  picked to match non-agricultural productivity difference of four. The aggregate productivity difference is the ratio of GDP per worker between the ninetieth and tenth percentile countries. Sector productivity differences are the ratios of sector output per worker in the ninetieth and tenth percentile countries. The ag/non-agriculture ratios are the agriculture productivity differences divided by the non-agriculture productivity differences.

*Source:* Authors' calculations and Caselli (2005).

selected based on agricultural productivity in both model countries, and hence average worker productivity is only slightly lower in the tenth-percentile country. Equation (8) from our analytical example illustrates this point: selection plays a larger role when employment shares by sector differ greatly, as they do in the main quantitative experiment, but not when the economies have similar sector employment shares, as in the current experiment.

#### B. Alternative Experiment No. 2: Calibrate to Non-agricultural Productivity Gaps and Agricultural Employment Shares

The second experiment chooses economy-wide efficiency differences to match the measured gap in non-agricultural productivity and, in addition, chooses an *agriculture-specific* efficiency difference to match the employment shares in agriculture in the ninetieth and tenth percentile countries.

To execute this experiment we introduce one additional parameter,  $A_a$ , which allows agricultural efficiency to differ from non-agricultural efficiency. Formally, our agricultural production function is now  $Y_a = A_a AL_a$ . While we take this additional parameter  $A_a$  as exogenous, it has several possible motivations. For one, it could represent agriculture-specific differences in land per worker or capital per worker, which we currently abstract from but explore in Section V. It could also represent the type of agriculture-specific distortion emphasized by the existing literature. For example, it could be distortions to the use of intermediate inputs in agriculture as studied by Restuccia, Yang, and Zhu (2008), or restrictions on farm size, as emphasized in Adamopoulos and Restuccia (2011).

Beginning from the benchmark model calibrated as in the main experiment, we normalize  $A_a$  to be one in the United States. We then lower  $A$  and  $A_a$  to match a productivity difference of four in non-agriculture, as in the experiments of Restuccia, Yang, and Zhu (2008) and Adamopoulos and Restuccia (2011), and an agricultural employment share of 78 percent, as in the tenth percentile country. We then compute the model's predictions for aggregate GDP per worker and agricultural productivity in the ninetieth and tenth percentile countries.

Table 6 presents the results of the second alternate experiment. Note that the gap in non-agricultural productivity is a factor of four in the model (as in the data) by construction. This is generated from an underlying  $A$  difference of eight (and not

TABLE 6—90–10 PRODUCTIVITY DIFFERENCES, ALTERNATIVE EXPERIMENT 2

	Agriculture	Aggregate	Non-agriculture	Ag/non-agriculture ratio
Data	45	22	4	10.7
Model	41	16	4	10.0
Without selection	27	17	8	3.4

*Notes:* This table displays results with  $A$  and  $A_a$  picked to match non-agricultural productivity ratio of four and a 78 percent agricultural employment share in the poor country. The aggregate productivity difference is the ratio of GDP per worker between the ninetieth and tenth percentile countries. Sector productivity differences are the ratio of sector output per worker in the ninetieth and tenth percentile countries. The ag/non-agriculture ratios are the agriculture productivity differences divided by the non-agriculture productivity differences.

*Source:* Authors' calculations and Caselli (2005).

four). The reason is that the model endogenously leads to lower non-agriculture differences than  $A$  differences, as explained above.<sup>15</sup> The difference in aggregate productivity is a factor 16, modestly less than in the data, and the difference in agricultural productivity is now a factor of 41, not far below the factor 45 in the data. Overall, the model delivers productivity differences that are ten times as large in agriculture as non-agriculture, just slightly less than the data.

To measure the importance of selection, we re-solve the model without the selection mechanism. In this case the model predicts a slightly larger aggregate difference of 17, a smaller agricultural difference of 27, and a non-agricultural difference of eight (which is exactly the  $A$  difference). Agricultural productivity is now just 3.4 times as variable across countries as non-agricultural productivity (which is exactly the  $A_a$  difference.) Thus, the importance of the selection is given by the ratio of sector productivities in the model with selection divided by the ratio in the model without selection, which is 2.9 (10.0/3.4). This is comparable to the value of 2.2 resulting from the main experiment.

## V. Extended Model with Capital and Land

We now extend the model to include capital and land. Up to this point we abstracted from capital and land mainly for transparency. One concern with this abstraction is that capital and land may interact with selection in ways that diminish the importance of selection. A second concern is that, by ignoring capital and land, the calibration procedure may overestimate the amount of wage variation that is attributable to productivity variation across individuals, which would again lead to an overestimate of the importance of selection. As we show below, neither of these concerns turn out to be warranted.

<sup>15</sup>In contrast, the experiments of Restuccia, Yang, and Zhu (2008) and Adamopoulos and Restuccia (2011) require  $A$  differences of the exact same size as the non-agricultural productivity differences.

### A. Environment

In this extended model, each worker has access to technologies to produce either the agricultural good or the non-agricultural good. The technologies are

$$y_a^i = Ak^{\phi_k} \ell^{\phi_l} (z_a^i)^{1-\phi} \quad \text{and} \quad y_n^i = Ak^\alpha (z_n^i)^{1-\alpha},$$

where  $\phi \equiv \phi_k + \phi_l$ ,  $k$  represents capital, and  $\ell$  represents land. Note that we abstract from land as a factor of production in the non-agricultural sector and allow for capital and labor’s shares to potentially differ across sectors, as consistent with recent estimates (Valentiny and Herrendorf 2008).

To solve this model we work backwards by first characterizing the solution to the profit maximization problem given an occupational choice, and then characterizing the occupational choice. Given the decision to work in agriculture, the profit maximization problem is

$$\max_{k, \ell} \{p_a Ak^{\phi_k} \ell^{\phi_l} (z_a^i)^{1-\phi} - rk - p_\ell \ell\},$$

where the price  $r$  is the cost of renting one unit of capital, and  $p_\ell$  is the price of renting one unit of land. Given the decision to work in non-agriculture, the profit maximization problem is

$$\max_k \{Ak^\alpha (z_n^i)^{1-\alpha} - rk\}.$$

Workers are the residual claimants on earnings after any payments to capital and land are made; we denote individual  $i$ ’s earnings as  $w^i(z_a^i)$  and  $w^i(z_n^i)$ , depending on which sector she chooses. Occupational choice comes down to a comparison of potential earnings in both sectors. These earnings are

$$(8) \quad w^i(z_a^i) = z_a^i (1 - \phi) (p_a A)^{\frac{1}{1-\phi}} \left(\frac{\phi_k}{r}\right)^{\frac{\phi_k}{1-\phi}}$$

$$\text{and} \quad w^i(z_n^i) = z_n^i (1 - \alpha) (A)^{\frac{1}{1-\alpha}} \left(\frac{\alpha}{r}\right)^{\frac{\alpha}{1-\alpha}}.$$

Combining the above equations yields a simple cutoff rule in relative individual productivity characterizing the optimal occupational choice for each worker. Working in non-agriculture is optimal for worker  $i$  if and only if

$$\frac{z_n^i}{z_a^i} \geq \chi p_a^{\frac{1}{1-\phi}} A^{\left(\frac{1}{1-\phi} - \frac{1}{1-\alpha}\right)} r^{\left(\frac{\alpha}{1-\alpha} - \frac{\phi_k}{1-\phi}\right)} p_\ell^{\frac{-\phi_l}{1-\phi}},$$

where  $\chi$  is a collection of constants. While similar to the cutoff rule in equation (3) of the benchmark model, this cutoff rule differs in two respects. First, the price of capital and the price of land now factor into the decision where to work. Second, economy-wide efficiency directly enters into the equation, with its impact determined by the difference in labor shares between the two sectors (i.e.,  $1 - \phi$  versus  $1 - \alpha$ ).

Optimal consumption decisions are the same as in the benchmark model. A worker's income now consists of her labor earnings plus an equal share of the aggregate payments to capital and land. An equilibrium of the economy consists of an agricultural price,  $p_a$ , a price of capital,  $r$ , a price of land,  $p_l$ , wages per efficiency unit of labor in each sector,  $w_a$  and  $w_n$ , and allocations for each worker, such that workers optimize and all markets clear.

### B. Calibration

We calibrate the preference parameters,  $\bar{a}$  and  $\nu$ , as well as the individual productivity parameters,  $\theta_a$ ,  $\theta_n$ , and  $\rho$ , as in the benchmark economy. Preferences do not change since only the production side of the model is different in the extended model. The reason the individual productivity distribution is calibrated the same is that in both models, log wage variation reflects only variation in log individual productivity, i.e.,  $\text{var}(\log w^i(z_a^i)) = \text{var}(\log z_a^i)$ . To see this in the extended model, equations in (8) show that payments to labor are proportional to individual productivity, and the degree of proportionality is common across workers of different productivity. Hence, calibrating the model to the same wage variance targets (and average wage target) described in Section IIIB results in the same three parameter values as in the benchmark model.

Incorporating capital and land adds several new parameters to calibrate: capital and land shares in agricultural production,  $\phi_k$  and  $\phi_l$ , capital's share in non-agricultural production,  $\alpha$ , and aggregate capital and land stocks, which we denote  $K$  and  $\mathcal{L}$ . We use the evidence of Valentinyi and Herrendorf (2008) on capital and land shares by sector in the United States to calibrate  $\phi_k$ ,  $\phi_l$ , and  $\alpha$ . They find values for capital and land's share in agricultural production to be 0.36 and 0.18 which we assign  $\phi_k$  and  $\phi_l$  to take. While these values are for the United States, they are also consistent with observed sharecropping arrangements in poor countries, where workers typically earn around one-half of all output; see Gollin, Lagakos, and Waugh (2012) for a more detailed discussion. For non-agriculture, Valentinyi and Herrendorf (2008) find capital's share to be 0.33, which we assign  $\alpha$  to take.

To calibrate the aggregate capital stock,  $K$ , we pick this value so the capital-output ratio in the rich economy is 2.5, which is consistent with evidence from the United States. To calibrate the aggregate land endowment we follow Adamopoulos and Restuccia (2011) and pick units such that average land per worker equals 169.3 hectares as they find in the United States data.

### C. Results

To explore the quantitative implications of the extended model, we normalize  $A$  to equal one and choose  $K$  to match GDP per worker relative to the United States for a country in the ninetieth percentile of the income distribution and a capital-output ratio of 2.5. We then lower  $A$  and  $K$  to match the aggregate productivity difference of 22 between the ninetieth and tenth percentile countries, and a capital-output ratio of one in the tenth percentile county. The latter is consistent with evidence from Caselli (2005) who computes capital-output ratios for the bottom 10 percent of countries to be approximately one.



TABLE 7—90–10 PRODUCTIVITY RATIOS, DATA AND EXTENDED MODEL

	Agriculture	Aggregate	Non-agriculture	Ag/non-agriculture ratio
Data	45	22	4	10.7
Model	38	22	10	3.8
Without selection	33	21	16	2.0
Without selection, land	17	15	14	1.1

*Notes:* The row “Without selection” reports productivity differences from the model with worker heterogeneity shut down. The row “Without selection and land” reports productivity differences from the model removing land as a factor of production. The aggregate productivity difference is the ratio of GDP per worker between the ninetieth and tenth percentile countries. Sector productivity differences are the ratio of sector output per worker in the ninetieth and tenth percentile countries. The ag/non-agriculture ratios are the agriculture productivity differences divided by the non-agriculture productivity differences.

*Source:* Authors’ calculations and Caselli (2005).

Table 7 presents the results. The extended model generates a factor 38 difference in productivity in agriculture and a factor ten difference in non-agriculture. This amounts to 3.8 times as much variation in agricultural productivity relative to non-agricultural productivity across countries. This is higher than the 2.2 ratio found in the benchmark experiment. Of course, the extended model has several other factors contributing to the larger differences in agriculture. In particular, land per worker is lower in the poor country, a feature that is present in other models with land as a fixed factor, such as Restuccia, Yang, and Zhu (2008), Adamopoulos and Restuccia (2011), and Herrendorf and Teixeira (2011).<sup>16</sup>

To isolate the importance of selection, we recompute the model’s predictions without the selection channel (i.e., with worker heterogeneity shut down) leaving all else the same. The third row in Table 7 reports the results. In this case the model predicts a lower agricultural difference of a factor 33, an aggregate difference of 21, and a higher non-agricultural difference of 16. The ratio of agriculture to non-agriculture productivity differences is 2.0. Thus, the model with selection generates 1.9 times larger variation in sectoral productivity than the model without selection (3.8/2.0), or not far below the 2.2 of the main experiment.

As a frame of reference, consider the importance of land relative to selection. The fourth row in Table 7 reports the results when we remove land from the model. Agricultural productivity differences now fall to a factor 17, aggregate differences fall to a factor 15, and non-agricultural differences fall to a factor 14. The ratio of agricultural to non-agricultural differences falls to 1.1. Using the same logic as above, land-per-worker differences contribute a factor 1.8 (2.0/1.1) to understanding the ratio of agricultural to non-agricultural productivity differences, which is of roughly similar magnitude to the selection channel.

<sup>16</sup>We find that the implications of this calibration for other cross-country observables not directly targeted, such as the share of labor in agriculture and relative price of agricultural goods, are reasonable. One additional check of the model is in the average size of a farm in the tenth percentile country. Our model predicts an average size of 4.6 hectares, which is quite close to the value of 5.0 hectares reported by Adamopoulos and Restuccia (2011).

An alternative way to measure the importance of selection is to consider the following decomposition of equilibrium output per worker in each sector. One can show that labor productivity in equilibrium can be written as

$$(9) \quad \frac{Y_a}{N_a} = (A)^{\frac{1}{1-\phi}} \left( \frac{K_a}{Y_a} \right)^{\frac{\phi_k}{1-\phi}} \left( \frac{\mathcal{L}}{Y_a} \right)^{\frac{\phi_l}{1-\phi}} \left( \frac{1}{N_a} \int_{i \in \Omega^a} z_a^i dGi \right), \quad \text{and}$$

$$(10) \quad \frac{Y_n}{N_n} = A^{\frac{1}{1-\alpha}} \left( \frac{K_n}{Y_n} \right)^{\frac{\alpha}{1-\alpha}} \left( \frac{1}{N_n} \int_{i \in \Omega^n} z_n^i dGi \right),$$

where the last bracketed terms in equations (9) and (10) represent the contribution from selection. Expressing output in this way has the benefit of giving “credit” for variations in  $K$  and  $\mathcal{L}$  generated by selection and differences in  $A$ . For example, agents with higher individual productivity optimally use more capital and land per unit of labor—but capital-output ratios and land-output ratios reflect only aggregate scarcity of  $K$  and  $\mathcal{L}$ . Klenow and Rodríguez-Clare (1997) and Hall and Jones (1999) make a similar argument for working with capital-output ratios rather than capital-labor ratios in the neoclassical growth model.

Taking a simple ratio of these bracketed terms in equations (9) and (10) across the rich and poor country decomposes the importance of each factor in accounting for the sector productivity differences in the model. The selection term contributes a factor of 1.55 to agriculture differences and 0.58 to non-agriculture differences. This implies that selection forces lead productivity differences in agriculture to be 2.7 times as large as those in non-agriculture (1.55/0.58), which is somewhat larger than in the main experiment.<sup>17</sup>

Together, these two decompositions establish bounds on the importance of selection in the extended model. The first decomposition suggests that selection leads to agriculture productivity differences that are 1.9 times as large as those of non-agriculture. The second decomposition suggests that selection leads to agriculture differences that are 2.7 times as large. Taken together, we conclude that the quantitative importance of selection is comparable in the extended model and benchmark model, and that land as a fixed factor and selection are complementary mechanisms.

## VI. Evidence: The Prevalence of Women in Agriculture across Countries

According to our theory, part of the large cross-country productivity differences in agriculture stem from poor countries having relatively more workers in agriculture who are unproductive at agricultural work (e.g., Table 3). In this section, we provide one concrete example of this phenomenon. In particular, we cite evidence that women are less productive at agricultural work than men on average, and we

<sup>17</sup>This decomposition suggests a limited role for land. Because land is fixed and agriculture output is lower in the poor country than in the rich country, land-to-output ratios are actually slightly higher in the poor country than in the rich country. This suggests that land plays no role in explaining agriculture productivity differences other than through the selection channel. Put differently, the fixed quantity of land is not a limiting factor for agriculture in poor countries given the low average productivity of their agriculture workers.

show that in cross-country data, women form a larger fraction of agricultural workers in developing countries than in richer countries.

A large body of literature has found that women tend to earn lower wages than men in agricultural work (e.g., Rosenzweig and Evenson 1977; Rosenzweig 1978; Psacharopoulos and Tzannatos 1992; and Horton 1996).<sup>18</sup> One widely proposed hypothesis for this gender wage gap in agriculture is that women are less productive at agricultural work than men on average (e.g., Goldin and Sokoloff 1982, 1984; Foster and Rosenzweig 1996; Pitt, Rosenzweig, and Hassan (2010), and Alesina, Giuliano, and Nunn 2011).

Several types of evidence support the hypothesis that women are less productive than men at agricultural work on average. As one piece of direct support, Pitt, Rosenzweig, and Hassan (2010) cite evidence from the United States and Bangladesh that men are physically stronger than women as measured by their grip strength. In Bangladesh, for example, 40 percent of men in a random sample of adults had a stronger grip than the strongest woman. This matters for productivity since much of agricultural work, such as plowing, is strength-intensive. Further support comes from the sexual division of labor in agriculture. Foster and Rosenzweig (1996) show that in the agricultural sectors of many developing countries, most men are hired to do plowing, while most women are hired to do weeding.<sup>19</sup> Goldin and Sokoloff (1982, 1984) argue that a major reason women earned less than men in agriculture in the early United States was that women were generally less productive at plowing than men.<sup>20</sup>

Given evidence that women are relatively less productive in agriculture than men, we next show that women constitute a relatively larger fraction of the agricultural work force in developing countries. In order to measure the prevalence of women in agriculture across countries, we draw on two independent sources of data. First, we use FAO data on the composition of agricultural workers by sex in 162 countries. The estimates come from a mix of labor force surveys and censuses of population. Second, we use data from the Integrated Public Use Microdata Series (IPUMS) provided by the Minnesota Population Center (2011) to compute the composition of agricultural workers by sex for 51 countries. These data come exclusively from nationally representative censuses of population, which in general have very large sample sizes. Using each dataset we compute the fraction of each country's agricultural workers who are women.

Figure 4 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 agricultural workers that are women.<sup>21</sup> For the countries with 70 percent or more of workers in

<sup>18</sup> Rosenzweig and Evenson (1977) and Rosenzweig (1978) document that, in India, women earn roughly 0.75 as much as men in agricultural work. Psacharopoulos and Tzannatos (1992) document gender wage gaps in agriculture of 0.92 in Colombia, 0.70 in Costa Rica, 0.76 in Guatemala, and 0.69 in Peru, and Horton (1996) documents gaps of 0.89 in Thailand and 0.85 in the Philippines.

<sup>19</sup> Alesina, Giuliano, and Nunn (2011) argue that, because women and children are less productive at plowing than men, societies that adopted plow agriculture earlier had lower demands for female and child labor and, hence, have lower fertility rates today.

<sup>20</sup> Goldin and Sokoloff (1984) document a larger gender wage gap in agriculture in the North than the South, and attribute it to the North's predominance of hay and wheat farming, where plowing was required, compared to the South's focus on tobacco and cotton, for which a smaller stature was useful in harvesting.

<sup>21</sup> Time series evidence from the development experiences of the United States and Britain paint a picture consistent with our cross-country evidence. Goldin and Sokoloff (1982, 1984) show that as the United States grew in

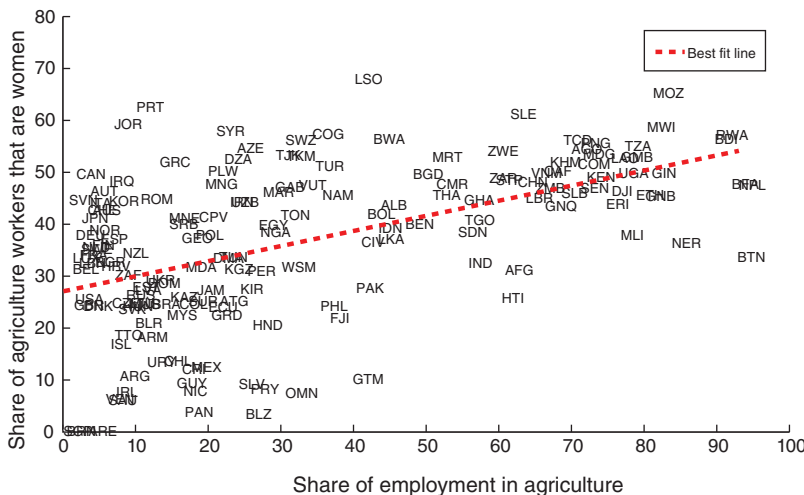


FIGURE 4. SHARE OF AGRICULTURAL WORKERS THAT ARE WOMEN

agriculture, roughly 50 percent of agricultural workers are women. In contrast, in countries with 10 percent of workers in agriculture or less, on average 30 percent of agricultural workers are women. A linear regression of the share of agricultural workers that are women on the share of all workers in agriculture yields a slope coefficient of 0.29 with a *p*-value of 0.01. The IPUMS data (not pictured) paints a similar picture: a similar linear regression using the IPUMS data yields a slope coefficient of 0.33 with a *p*-value of 0.01.<sup>22</sup>

Putting these pieces together—(i) women are the less productive at agricultural work and (ii) women are more prevalent in agriculture in developing countries—provides a concrete example of how agricultural productivity differences across countries depend on the average productivity of workers in the agricultural sector, as predicted by our theory.

### VII. Conclusion

We argue that cross-country productivity differences are larger in agriculture than in non-agriculture in part because of cross-country differences in the selection of heterogenous workers by sector. In poor countries, where economy-wide efficiency is low, subsistence food requirements lead workers that are relatively unproductive in agricultural work to nonetheless select into the agriculture sector. In rich countries, in contrast, those few workers self-selecting into agriculture are those who are

the nineteenth century, women shifted out of agriculture and into manufacturing much more rapidly than men. In Britain, the evidence of Allen (1994) shows that in 1700, 46 percent of adult agricultural workers were women, and by 1850 this fraction had fallen to just 29 percent (see Allen 1994, Table 5.3).

<sup>22</sup> One alternative theory for why women are more prevalent in agriculture in developing countries is that higher fertility rates in the developing world make work on the family farm—where childcare can be provided easily—particularly attractive for women. One piece of evidence against this alternative theory is that the share of women *without children* in agriculture also increases sharply in the agricultural share of employment. A linear regression using our IPUMS data of the share of agricultural workers that are female without children under five on the agricultural share of employment yields a slope coefficient of 0.21 with a *p*-value of 0.01.

relatively most productive at farm work. As a result, measured labor productivity differences are larger in agriculture than in the aggregate. Selection forces work in exactly the opposite way in non-agriculture, and productivity differences are smaller than those of the aggregate.

Quantitatively, we find that the selection channel leads agricultural productivity differences to be around twice as large as those of the non-agricultural sector. This result was found both in isolation and in the presence of competing mechanisms such as exogenous sector-specific productivity differences and capital and land differences. The key challenge our model faces is that for selection to work, it still requires large, exogenous productivity differences—either of a general or agriculture-specific nature—to draw workers into agriculture. Of course, what explains these differences is still an open question that both better measurement and theory can hopefully address.

## REFERENCES

- Acemoglu, Daron, Simon Johnson, and James A. Robinson.** 2001. “The Colonial Origins of Comparative Development: An Empirical Investigation.” *American Economic Review* 91 (5): 1369–401.
- Acemoglu, Daron, Simon Johnson, and James A. Robinson.** 2002. “Reversal of Fortune: Geography and Institutions in the Making of the Modern World Income Distribution.” *Quarterly Journal of Economics* 117 (4): 1231–94.
- Adamopoulos, Tasso.** 2009. “Do Richer Countries Have Higher Distribution Margins?” Unpublished.
- Adamopoulos, Tasso, and Diego Restuccia.** 2011. “The Size Distribution of Farms and International Productivity Differences.” Unpublished.
- Alesina, Alberto, Paola Giuliano, and Nathan Nunn.** 2011. “Fertility and the Plough.” *American Economic Review* 101 (3): 499–503.
- Allen, Robert.** 1994. “Agriculture During the Industrial Revolution.” In *The Economic History of Britain Since 1700, Volume 1*, edited by Roderick Floud and Deirdre McCloskey, 96–122. New York: Cambridge University Press.
- Atkeson, Andrew, and Masao Ogaki.** 1996. “Wealth-Varying Intertemporal Elasticities of Substitution: Evidence from Panel and Aggregate Data.” *Journal of Monetary Economics* 38 (3): 507–34.
- Beegle, Kathleen, Joachim De Weerd, and Stefan Dercon.** 2011. “Migration and Economic Mobility in Tanzania: Evidence from a Tracking Survey.” *Review of Economics and Statistics* 93 (3): 1010–33.
- Butzer, Rita, Yair Mundlak, and Donald F. Larson.** 2010. “Measures of Fixed Capital in Agriculture.” The World Bank, Policy Research Working Paper 5472.
- Case, Anne, and Christina Paxson.** 2008. “Stature and Status: Height, Ability, and Labor Market Outcomes.” *Journal of Political Economy* 116 (3): 499–532.
- Caselli, Francesco.** 2005. “Accounting for Cross-Country Income Differences.” In *Handbook of Economic Growth*, edited by Philippe Aghion and Steven Durlauf, 679–741. Amsterdam: Elsevier Science.
- Caselli, Francesco, and Wilbur John Coleman, II.** 2001. “The US Structural Transformation and Regional Convergence: A Reinterpretation.” *Journal of Political Economy* 109 (3): 584–616.
- Center for Disease Control.** 2009. “National Health Interview Survey.” <http://www.cdc.gov/nchs/nhis.htm> (accessed August 1, 2011).
- Center for Economic Policy Research.** 1996–2010. “Current Population Survey, March. CEPR Uniform Data Extracts.” <http://ceprdata.org/> (accessed August 1, 2011).
- Chanda, Areendam, and Carl-Johan Dalgaard.** 2008. “Dual Economies and International Total Factor Productivity Differences: Channelling the Impact from Institutions, Trade, and Geography.” *Economica* 75 (300): 629–61.
- Donovan, Kevin.** 2011. “Agricultural Risk, Intermediate Inputs, and Cross-Country Productivity Differences.” Unpublished.
- Eaton, Jonathan, and Samuel Kortum.** 2002. “Technology, Geography, and Trade.” *Econometrica* 70 (5): 1741–79.
- Food and Agriculture Organization.** 2004. “FAO Statistical Yearbook; Table A.3 Economically active population in agriculture.” <http://www.fao.org/docrep/008/y5473m/y5473m00.HTM> (accessed August 1, 2011).



- Food and Agriculture Organization.** 2004. "FAO Statistical Yearbook; Table G.2 Agricultural GDP and its share in total GDP." <http://www.fao.org/docrep/008/y5473m/y5473m00.HTM> (accessed August 1, 2011).
- Foster, Andrew D., and Mark R. Rosenzweig.** 1996. "Comparative Advantage, Information and the Allocation of Workers to Tasks: Evidence from an Agricultural Labour Market." *Review of Economic Studies* 63 (3): 347–74.
- Goldin, Claudia, and Kenneth Sokoloff.** 1982. "Women, Children, and Industrialization in the Early Republic: Evidence from the Manufacturing Censuses." *Journal of Economic History* 42 (4): 741–74.
- Goldin, Claudia, and Kenneth Sokoloff.** 1984. "The Relative Productivity Hypothesis of Industrialization: The American Case, 1820 to 1850." *Quarterly Journal of Economics* 99 (3): 461–87.
- Gollin, Douglas, David Lagakos, and Michael E. Waugh.** 2011. "Selection, Agriculture, and the Gains from Trade." Unpublished.
- Gollin, Douglas, David Lagakos, and Michael E. Waugh.** 2012. "The Agricultural Productivity Gap in Developing Countries." Unpublished.
- Hall, Robert E., and Charles I. Jones.** 1999. "Why Do Some Countries Produce So Much More Output Per Worker Than Others?" *Quarterly Journal of Economics* 114 (1): 83–116.
- Heckman, James J., and Bo E. Honore.** 1990. "The Empirical Content of the Roy Model." *Econometrica* 58 (5): 1121–49.
- Heckman, James J., and Guilherme Sedlacek.** 1985. "Heterogeneity, Aggregation, and Market Wage Functions: An Empirical Model of Self-selection in the Labor Market." *Journal of Political Economy* 93 (6): 1077–125.
- Herrendorf, Berthold, and Todd Schoellman.** 2011. "Why is Measured Productivity so Low in Agriculture?" Unpublished.
- Herrendorf, Berthold, and Arilton Teixeira.** 2011. "Barriers to Entry and Development." *International Economic Review* 52 (2): 573–602.
- Herrendorf, Berthold, and Akos Valentinyi.** 2012. "Which Sectors Make Poor Countries So Unproductive?" *Journal of the European Economic Association* 10 (2): 323–41.
- Heston, Alan, Robert Summers, and Bettina Aten.** 2011. "Penn World Table Version 7.0." Center for International Comparisons of Production, Income, and Prices at the University of Pennsylvania, June 2011. <https://pwt.sas.upenn.edu/icp.html> (accessed August 1, 2011).
- Horton, Susan, ed.** 1996. *Women and Industrialization in Asia*. New York: Routledge.
- Hsieh, Chang-Tai, and Peter J. Klenow.** 2007. "Relative Prices and Relative Prosperity." *American Economic Review* 97 (3): 562–85.
- Jones, Charles I.** 1994. "Economic Growth and the Relative Price of Capital." *Journal of Monetary Economics* 34 (3): 359–82.
- Klenow, Peter J., and Andres Rodríguez-Clare.** 1997. "The Neoclassical Revival in Growth Economics: Has It Gone Too Far?" In *National Bureau of Economic Research Macroeconomics Annual 1997*, edited by Ben S. Bernanke and Julio J. Rotemberg, 73–103. Cambridge, MA: MIT Press.
- Lagakos, David, and Michael E. Waugh.** 2013. "Selection, Agriculture, and Cross-Country Productivity Differences: Dataset." *American Economic Review*. <http://dx.doi.org/10.1257/aer.103.2.948>.
- Lanzona, Leonardo A.** 1998. "Migration, Self-Selection and Earnings in Philippine Rural Communities." *Journal of Development Economics* 56 (1): 27–50.
- Miguel, Edward, and Joan Hamory.** 2009. "Individual Ability and Selection into Migration in Kenya." United Nations Development Programme, Human Development Research Paper 2009/45.
- Minnesota Population Center.** 2011. *Integrated Public Use Microdata Series, International: Version 6.1* [Machine-readable database]. Minneapolis: University of Minnesota.
- Nelsen, Roger B.** 2006. *An Introduction to Copulas*. second ed. New York: Springer.
- Pitt, Mark M., Mark R. Rosenzweig, and Md. Nazmul Hassan.** 1990. "Productivity, Health, and Inequality in the Intrahousehold Distribution of Food in Low-Income Countries." *American Economic Review* 80 (5): 1139–56.
- Pitt, Mark M., Mark R. Rosenzweig, and Nazmul Hassan.** 2012. "Human Capital Investment and the Gender Division of Labor in a Brawn-Based Economy." *American Economic Review* 102 (7): 3531–60.
- Prasada Rao, D.S.** 1993. *Intercountry Comparisons of Agricultural Output and Productivity*. Rome: Food and Agriculture Organization of the United Nations.
- Psacharopoulos, George, and Zafiris Tzannatos, eds.** 1992. *Case Studies on Women's Employment and Pay in Latin America*. Washington, DC: World Bank.
- Restuccia, Diego, and Carlos Urrutia.** 2001. "Relative Prices and Investment Rates." *Journal of Monetary Economics* 47 (1): 93–121.
- Restuccia, Diego, Dennis Tao Yang, and Xiaodong Zhu.** 2008. "Agriculture and Aggregate Productivity: A Quantitative Cross-Country Analysis." *Journal of Monetary Economics* 55 (2): 234–50.

- Rosenzweig, Mark R.** 1978. "Rural Wages, Labor Supply, and Land Reform: A Theoretical and Empirical Analysis." *American Economic Review* 68 (5): 847–61.
- Rosenzweig, Mark R., and Robert E. Evenson.** 1977. "Fertility, Schooling, and the Economic Contribution of Children in Rural India: An Econometric Analysis." *Econometrica* 45 (5): 1065–79.
- Rosenzweig, Mark R., and Kenneth I. Wolpin.** 1993. "Credit Market Constraints, Consumption Smoothing, and the Accumulation of Durable Production Assets in Low-Income Countries: Investment in Bullocks in India." *Journal of Political Economy* 101 (2): 223–44.
- Roy, A. D.** 1951. "Some Thoughts on the Distribution of Earnings." *Oxford Economic Papers New Series* 3: 135–46.
- Schultz, T.W.** 1953. *The Economic Organization of Agriculture*. New York: McGraw Hill.
- Steckel, Richard H.** 1995. "Stature and the Standard of Living." *Journal of Economic Literature* 33 (4): 1903–40.
- Strauss, John, and Duncan Thomas.** 1998. "Health, Nutrition, and Economic Development." *Journal of Economic Literature* 36 (2): 766–817.
- Summers, Robert, and Alan Heston.** 1991. "The Penn World Table (Mark 5): An Expanded Set of International Comparisons, 1950–1988." *Quarterly Journal of Economics* 106 (2): 327–68.
- Teignier, Marc.** 2012. "The Role of Trade in Structural Transformation." Unpublished.
- Tombe, Trevor.** 2011. "The Missing Food Problem: How Low Agricultural Imports Contribute to International Income and Productivity Differences." Unpublished.
- Valentinyi, Akos, and Berthold Herrendorf.** 2008. "Measuring Factor Income Shares at the Sectoral Level." *Review of Economic Dynamics* 11 (4): 820–35.
- Vollrath, Dietrich.** 2009. "How Important Are Dual Economy Effects for Aggregate Productivity?" *Journal of Development Economics* 88 (2): 325–34.
- World Bank.** 2005. "International Comparison Program." <http://data.worldbank.org/data-catalog/international-comparison-program> (accessed August 1, 2011).