

Wages, Human Capital, and the Allocation of Labor across Sectors^{*}

Berthold Herrendorf and Todd Schoellman[†]

March 7, 2014

Abstract

We document for nine countries ranging from rich (Canada, U.S.) to poor (India, Indonesia) that average wages are higher in non-agriculture than in agriculture. We measure sectoral human capital and find that it accounts for the entire wage gap in the U.S. and most of the wage gaps elsewhere. We develop a multi-sector model that explains these findings if: (i) Mincer returns to schooling are equal in both sectors; (ii) more able workers sort into non-agriculture; (iii) distortions to the allocation of labor between sectors are negligible in the U.S. and small elsewhere.

^{*}We would like to thank Gustavo Ventura for many helpful discussions. Herrendorf thanks the Spanish Ministry of Education for research support (Grant ECO2012-31358) and the University of Mannheim for its hospitality during a sabbatical when this project originated. The usual disclaimer applies.

[†]Address: Department of Economics, W.P. Carey School of Business, Arizona State University, Tempe, AZ 85287-9801, USA. E-mails: berthold.herrendorf@asu.edu and todd.schoellman@gmail.com

1 Introduction

In his Nobel prize lecture, Kuznets (1973) included the process of structural transformation as one of the six main features of modern economic growth. Structural transformation refers the reallocation of labor from agriculture into other sectors of the economy. Perhaps somewhat surprisingly there are two rather different views of structural transformation within the macro-development literature. One strand of the literature *assumes* that at each moment labor is allocated efficiently between sectors and derives the properties of preferences and technological progress that generate structural transformation as a consequence of growth; see e.g. Herrendorf, Rogerson and Valentinyi (2013) and the references therein. In contrast, a second strand of the literature measures large gaps in labor productivity between agriculture and the rest of the economy. Since it is difficult to account for these gaps, this strand of the literature concludes that labor must be allocated *inefficiently* between sectors, with the reallocation of labor in poor countries hindered by what is typically referred to as barriers or wedges; see e.g. Caselli (2005) and Restuccia, Yang and Zhu (2008). According to this logic, the removal of barriers would allow labor to move out of agriculture and would generate growth. This is precisely the opposite direction of causality as in the first strand of the literature. It follows that given the same stylized fact these two strands of the macro-development literature will make vastly different policy prescriptions.

Somewhat surprisingly, there is little hard evidence on whether or not the allocation of labor between non-agriculture and agriculture is efficient. The contribution of this paper is to provide such evidence for multiple years and nine countries ranging from rich (Canada, U.S.) to poor (India, Indonesia). The selection criterion for our countries is access to sufficiently detailed data that allow us to calculate wages and construct human capital at the sectoral level. For the U.S. these data are from the CPS, the Census, and the American Community Survey (ACS) and they cover the period 1980–2011. For the other eight countries, these data are from 26 population censuses that are harmonized by IPUMS and range from 1970 to 2010. We document that there are large gaps between the average wages per worker in non-agriculture and agriculture in all countries including the U.S. We decompose the wage gaps into differences in the observable characteristics of workers; differences in the returns to observable characteristics; residual wage gaps.

For the U.S., we find that the wage gaps are entirely accounted for by differences in observables: non-agricultural workers have more years of schooling and earn higher Mincer returns to schooling than agricultural workers. There two different interpretations of this finding. The “sectoral view” attributes the differences in Mincer returns to sectoral

technologies: schooling generates less human capital for workers who choose agriculture. The “selection view” attributes the differences to the workers: schooling generates the same human capital in each sector, but workers sort according to their unobserved ability. Both views have the potential to explain why non-agriculture has more human capital and higher Mincer returns than agriculture. To help us distinguish between them, we construct a multi-sector model in which workers choose in which sector they work. We assume that workers differ in their unobserved innate ability and their years of schooling and that sectors potentially differ in the technology that translates these characteristics into human capital. We also allow for the possibility of distorting barriers/wedges. We show that to account for the U.S. findings we need to restrict our model along two dimensions: (i) both sectors have the same technology; (ii) there are no barriers/wedges to the allocation of labor. Note that (i) implies that a given worker gets the same returns to his characteristics in both sectors, which of course is the selection view. Under these restrictions, our model has a sorting equilibrium with the following properties: workers in non-agriculture have more years of schooling, higher innate ability, and more human capital; average wages per *worker* are higher in non-agriculture; average wages per *efficiency unit* are equalized between the two sectors. The restriction that both sectors have the same technology has the testable implication that individuals who switch between sectors should not experience sizeable changes in the returns to their observable characteristics. We provide evidence from the CPS that is broadly consistent with that view.

For the other countries, we find that most of the wage gaps are accounted for by differences in observables: non-agricultural workers once again have more years of schooling and earn higher Mincer returns on schooling; differences in the residual wage gap are small in comparison to the differences in human capital. Imposing the selection view that the sectoral technologies are the same in both sectors, we show that our model implies that the residual wage gap provides an upper bound on the size of barriers/wedges. We conclude that barriers/wedges that lead to the mis-allocation of labor between the agriculture and non-agriculture account for a relatively small part of the raw wage gap between non-agriculture and agriculture.

Our work has several important implications for the macro-development literature. To begin with, it implies that to construct human capital at the sector level one does not only need to take into account selection according to observed characteristics but also selection according to unobserved innate ability. Mincer returns estimated at the sectoral level deliver this, whereas Mincer returns estimated at the aggregate level miss the selection according to unobserved innate ability. Second, our finding that in the U.S. wages per efficiency unit are

equalized between non-agriculture and agriculture is a necessary condition for the efficiency of the allocation of labor between the two sectors, which part of the literature on structural transformation assumes. However, it is not a sufficient condition because it does not rule out distortions that affect the allocation of labor without affecting wages, an example being production subsidies to agriculture. Lastly, our finding that even in the poorest countries of our sample the potential role for barriers/wedges is relatively small limits the scope for policy reforms that aim to generate growth by removing barriers/wedges.

The remainder of the paper is organized as follows. In the next section, we present our basic findings for the U.S. Section 3 views the basic findings through the lens of a simple multi-sector model. Section 4 presents further evidence for the U.S. Section 5 extends the U.S. analysis to a cross section of eight countries. Section 6 discusses the implications of our findings. Section 7 concludes and provides suggestions for future research.

2 Basic Facts About Sectoral Wages and Human Capital in the U.S.

2.1 Overview and Basics

Our goal in this section is to describe the stylized facts about wages in agriculture vs. non-agriculture. We start with the U.S. for which we have rich data on cross sections. In a later section, we will turn to international data and show that broadly similar patterns apply there.

Within the U.S. there are two data sets that are nationally representative, have useful data on wages, schooling, age, and so on, and have a sufficient sample size in agriculture to produce useful statistics: the population census and the Current Population Survey (CPS).¹ Within the CPS there are two widely used ways to construct data on wages. The first way is to use the monthly files. Workers in the outgoing rotation group (those in the fourth or eighth month in the sample) respond to a full battery of questions that we need, so we can use 1/4 of the sample every month. A second way is to use the March demographic oversample. For now we show results from all three data sets. They will agree on broad trends. We restrict our attention to the period 1980–2011. The main reason for this is that we lack some of the critical data (in particular on switchers) before 1980.

¹The 2000 Census was the last to include the “long form” that provides the detailed data on a sample of households. Since 2000 this information has been collected annually from a smaller sample of households through the American Community Survey (ACS). The questions and responses are quite similar, so we combine ACS data for year 2005–2011 with the Censuses of 1980, 1990, and 2000.

In all three data sets we focus on the subsample of workers that are typically used in wage regressions. We limit the sample to workers who have valid responses to the questions of interest (industry of employment, income, and so on). We use only workers 18–65 years old. For some of our results we will want to understand the role of potential experience, which we define as age minus schooling minus 6; we include in our sample only workers who have 0–45 years of potential experience. We use workers who are employed, work for wages, and report positive wage income for the relevant period. While these restrictions are all standard, we note that the restriction to wage workers is somewhat stronger here than is usually the case. The reason is that roughly half of the agricultural labor force is self-employed proprietors. While these individuals can be of interest, we avoid studying them for two reasons. First, their income represents payments to both capital and labor, and it is hard to disentangle the fraction of proprietors’ income that is wage income. Second, it is well known that these proprietors underreport their income by a large amount, in which case we want to avoid taking their stated income too seriously (Herrendorf and Schoellman 2013). The income of wage workers is much better reported and easier to interpret.

We classify workers into agriculture and non-agriculture on the basis of their reported industry of employment. We define agriculture as crop and livestock farming. All other industries are allocated to non-agriculture. The appendix includes details on the exact industry responses available by year and how they are allocated.

2.2 Stylized Fact 1: Large Raw Wage Gaps

We start by documenting the wage gap, defined as the ratio of the average hourly wages in agriculture to non-agriculture, both expressed in current dollars. A wage gap smaller than one indicates that average wages in agriculture are smaller than the average wages in non-agriculture. To find the wage gap, we estimate a simple wage regression:

$$\log(w_{ijt}) = \beta_d d_t + \beta_{dj} d_j + \varepsilon_{ijt}$$

where w_{ijt} is the hourly wage of individual i in sector j during year t , d_t is our parsimonious notation for a full set of year dummies, d_j is our notation for a sector dummy, β_d and β_{dj} are coefficients, and ε_{ijt} is an iid error term with zero mean. The point of this regression is to control for average wage growth and inflation by year through the full set of year dummies and then estimate the average sectoral wage gap. The log-wage gap between agriculture and non-agriculture is captured by β_{da} if non-agriculture is the omitted group. We report $\exp(\beta_{da})$, i.e. the ratio of average hourly wages in agriculture to non-agriculture.

Table 1: Raw Wage Gaps in the U.S. 1980–2011

	U.S. Census	March CPS	Monthly CPS
Raw Wage Gap	0.60	0.56	0.58
Adjusted Wage Gap	0.57	0.53	0.55

We estimate this regression separately for all three data sources. The wage we construct differs slightly by source. In the Census/ACS, it is constructed as last year’s income divided by the product of hours usually worked in a week times weeks worked in the year. The wage for the March CPS is constructed in the same way. The wage for the monthly CPS file is a bit different, as in the monthly CPS file workers report either their hourly wage directly or their weekly earnings and hours worked for the prior week.

Details aside, the picture is quite consistent across data sets, as shown in Table 1. Hourly wages in agriculture are only between 50–60% of those in non-agriculture. This is the essential fact we want to shed light on in this paper.

As a first pass, we consider whether this raw wage gap can be accounted for by controlling for some simple demographic factors. To this end, we estimate an adjusted wage gap, which equals the value of $\exp(\beta_{da})$ in the following regression:

$$\log(w_{ijt}) = \beta_d d_t + \beta_{dj} d_j + \beta_x X_{ijt} + \varepsilon_{ijt}$$

where X_{ijt} is the set of controls. We include state (because some states persistently have higher wages, perhaps as a compensating differential for high cost of living) and gender (because women are paid less on average than men). The adjusted wage gap $\exp(\beta_{da})$ is the same wage gap as before except that we have netted out the potentially confounding effects of state and gender. The last row in Table 1 gives the results and shows that they hardly differ. In fact, the adjusted wage gaps are larger than the unadjusted wage gaps. The reason for this is that agriculture has so many male workers whose wages are higher than those for women in the same sector. We now turn to understanding the source of these adjusted wage gaps.

2.3 Stylized Fact 2: Differences in Observed Characteristics

A natural candidate explanation for adjusted wage gaps is sectoral differences in human capital. After all, most existing theories of how workers choose their sector predict equalization of the wage per efficiency unit of labor, not equalization of the wage per hour worked. Hence,

Table 2: Gaps in Observed Characteristics (U.S. 1980–2011)

	U.S. Census	March CPS	Monthly CPS
Years of Schooling	3.3	3.1	3.2
Years of Potential Experience	0.0	0.4	0.5

we adjust wages for sectoral differences in human capital. The data also suggest that this adjustment may be important, because the observable characteristics typically associated with human capital turn out to differ across sectors. Table 2 shows this: non-agricultural workers have more than three years more schooling and half a year more experience than agricultural workers. Since the second difference is rather small, we abstract from it for now and focus our analysis on years of schooling.

To explore the quantitative role of schooling for explaining wages, we need to translate years of schooling into human capital. As a useful first step we follow the approach pioneered by Bils and Klenow (2000). They show that under simple mild assumptions, the log-human capital gain from an additional year of schooling is equal to the log-wage gain (or Mincer return) from an additional year of schooling. We consider two applications of this idea that vary only in the Mincer returns to schooling that we use to construct human capital stocks. We first work with the standard Mincer returns used in the development accounting literature and popularized by Hall and Jones (1999); we then work with the actual Mincer returns estimated from our data.

Using the actual estimated returns to schooling in our data turns out to be complicated slightly by the fact that the data are not well fit by a linear return to schooling. Instead, wages are a convex function of schooling in all three data sources that we use. This is consistent with recent work in the U.S. as well as in many other countries around the world (Lemieux 2006, Binelli 2012). To show the relationship in our data, we estimate regressions that are fully flexible in schooling:

$$\log(w_{ijt}) = \beta_d d_t + \beta_{di} d_i + \beta_{ds} d_s + \beta_x X_{ijt} + \varepsilon_{ijt}$$

where d_s is a full set of dummies for years of schooling. The results are plotted in Figure 1. There is a notable change in slope that arises at roughly twelve years of schooling in all three figures. This suggests the parsimonious regression that is used for the remainder of our results,

$$\log(w_{ijt}) = \beta_d d_t + \beta_{di} d_i + \beta_x X_{ijt} + \beta_s S_{ijt} + \beta_c C_{ijt} + \varepsilon_{ijt}$$

where S is years of schooling (primary and secondary education) and C is years of college (tertiary education, that is, years of schooling past the twelfth). This functional form allows for a linear return to schooling β_s in the standard way, but also captures through β_c that college has a higher return.²




Figure 1: Convexity in the Returns to Schooling

Table 3 shows the human capital gaps estimated using the Hall–Jones rate of returns and our estimated rate of returns. The results of the two procedures are fairly similar: both find that human capital in agriculture is about 80% of human capital in non-agriculture. The human capital gaps constructed in this way are smaller than the gaps in wages. An immediate implication is that they do not account for the entirety of the wage gaps. Table 4

²Note that most of the literature follows Psacharopoulos (1994) and assumes that wages are a concave function of schooling, at odds with the current data. We show below that this has little effect on our results.

Table 3: Standard Estimates of Human Capital Gaps (U.S. 1980–2011)

	U.S. Census	March CPS	Monthly CPS
Hall–Jones Rate of Return	0.80	0.81	0.80
Estimated Rate of Return	0.80	0.79	0.80

Table 4: Wage Gaps Adjusted for Standard Human Capital Gaps (U.S. 1980–2011)

	U.S. Census	March CPS	Monthly CPS
Hall–Jones Rates of Return	0.72	0.70	0.69
Estimated Rates of Return	0.72	0.67	0.69

makes this point precise. The line entitled “estimated rate of return” shows the unaccounted for wage gaps, that is, the remaining gaps in the level of wages between agriculture and non–agriculture after netting out the effects of schooling. Agriculture is still estimated to offer roughly 30% lower wages than non–agriculture. To understand why, we now develop a simple model of sectoral wages and human capital which is going to guide the next steps of our empirical analysis.

3 Model

3.1 Environment

Consider a static environment with $L > 0$ individuals. Each individual is endowed with one unit of time, an innate ability x and non–negative numbers of years of schooling $s \geq 0$. Denote by Ω the two–dimensional set of individual characteristics.

There is an agricultural and non–agricultural consumption good. Individual preference over the consumption of the two goods are represented by the following utility function:

$$u(c^a, c^n) = \alpha \log(c^a) + (1 - \alpha) \log(c^n) \quad (1)$$

where the superscripts indicate the sector and $\alpha \in (0, 1)$ is a relative weight. For simplicity, we will assume that individuals do not value leisure, implying that in equilibrium they will allocate their full time endowment to work.

The two goods are produced according to linear production functions that use labor as

the only input factor:

$$Y^i = A^i H^i \quad (2)$$

A^i denotes TFP, which here is the same as labor productivity, and H^i denotes human capital in sector i .³

$$H^i = \int_{\Omega} \nu^i(x, s) h^i(x, s) d\mu(x, s) \quad (3)$$

where $\nu^i(x, s)$ is an indicator function that equals one if individual (x, s) works in sector i . $h^i(x, s)$ is the human capital of individual (x, s) in sector i . For future reference, we denote the number of individuals in sector i by L^i . Note that when we estimated H^i above, we imposed that the returns to schooling are the same in both sectors. This is the case if $h^a(x, s) = h^n(x, s)$. In general, however, there is no reason to restrict $h^i(x, s)$ to be the same in both sectors, so we don't impose this on the model from the start.

The development literature considers barriers/wedges that prevent the free movement of labor from agriculture to non-agriculture and distort the efficient allocation of labor between the two sectors. A simple way of capturing the effects of such barriers on the allocation of labor is to assume that there is a government that imposes a tax τ on wages in non-agriculture and redistributes the proceeds through a lump-sum transfer T to all individuals. This is similar to the way in Restuccia et al. (2008) modeled barriers.

3.2 Equilibrium definition

A competitive equilibrium in this model world is goods prices (P^a, P^n) (one of which is set to one reflecting the choice of numeraire), rental rates (W^a, W^n) , a tax rate τ , choices of consumption and sector $\{(c^a, c^n, \nu^a, \nu^n)(x, s)\}_{(x,s) \in \Omega}$, output and labor in each sector Y^i, H^i such that:

- $\forall (x, s) \in \Omega$, $(c^a, c^n, \nu^a, \nu^n)(x, s)$ solve the individual problem:

$$\begin{aligned} \max_{c^a, c^n, \nu^a, \nu^n} \quad & \alpha \log(c^a) + (1 - \alpha) \log(c^n) \\ \text{s.t.} \quad & P^a c^a + P^n c^n = W^a h^a(x, s) \nu^a + (1 - \tau) W^n h^n(x, s) \nu^n + T \end{aligned} \quad (4)$$

³Note that H^i is sometimes referred to as efficiency units in sector i .

- (Y^i, H^i) solve the firm problem in sector i :

$$\max_{Y, H} P^i Y - W^i H \quad \text{s.t.} \quad Y = A^i H \quad (5)$$

- Markets clear:

$$Y^i = \int_{\Omega} c^i(x, s) d\mu(x, s) \quad (6)$$

$$H^i = \int_{\Omega} \nu^i(x, s) h^i(x, s) d\mu(x, s) \quad (7)$$

One might be tempted to start solving this model by requiring that the rental rates of human capital are equalized across sectors, $W^a = W^n$. For two reasons this is not necessarily an equilibrium property of our simple model. The first one is that the W^i 's have to adjust in such a way that the labor market clears and the “right” number of individuals chooses each sector. The second reason is that we have assumed that only W^n is taxed, implying that $W^n > W^a$ even if the net wages were equalized across sectors, $(1 - \tau)W^n = W^a$.

3.3 Equilibrium sorting

Since in the data individuals in non-agriculture have more years of schooling, we are interested in equilibria in which individuals sort according to their characteristics. To ensure that such an equilibrium exists and is unique in our model, we need to impose more structure on the environment. It is natural to assume that schooling is a continuous and increasing function of the innate ability:

$$s(x), s'(x) \geq 0 \quad (8)$$

This would result endogenously in a dynamic version of this model where individuals optimally choose schooling. It is also natural to assume that human capital is a continuous and increasing function of both schooling and ability. Together these two assumptions imply that:

$$\frac{dh^i(x, s(x))}{dx} > 0 \quad (9)$$

where dh^i/dx is the total derivative of h^i with respect to x .

If positive measures of individuals work in each sector, then there must at least be one individual x_l who weakly prefers agriculture and one individual x_h who prefers non-agriculture:

$$W^a h^a(x^l, s(x^l)) \geq (1 - \tau) W^n h^n(x^l, s(x^l)) \quad (10)$$

$$W^a h^a(x^h, s(x^h)) \leq (1 - \tau) W^n h^n(x^h, s(x^h)) \quad (11)$$

Continuity then implies that there must be an indifferent individual $(x^*, s(x^*))$ with:

$$W^a h^a(x^*, s(x^*)) = (1 - \tau) W^n h^n(x^*, s(x^*))$$

Note that if

$$h^a(x^*, s(x^*)) \neq h^n(x^*, s(x^*))$$

then

$$W^a \neq (1 - \tau) W^n$$

that is, the rental rates of efficiency units may differ in our model even if there is no tax distortion. The reason for this, of course, is that if the technologies that translate individual characteristics into human capital differ across sectors, then the net wage per efficiency unit needs to be higher in the sector in which individual $(x^*, s(x^*))$'s human capital is lower. This is important to keep in mind, particularly since some contributors to the literature act as if efficiency (here $\tau = 0$) was the same as the equalization of the rental rates of efficiency units across sectors.

Given the evidence collected in the previous section, it is natural to focus on a sorting equilibrium in which all individuals with $x < x^*$ choose agriculture and all individuals with $x > x^*$ choose non-agriculture. A sufficient conditions for the uniqueness of such a sorting equilibrium is that at each intersection wages in agriculture changes less in ability than wages in non-agriculture:

$$\frac{dW^a h^a(x^*, s(x^*))}{dx} < \frac{dW^n h^n(x^*, s(x^*))}{dx}$$

This implies if the two wages intersect, then the wage in non-agriculture cuts the wage in agriculture from below.

We are now in a position to state several properties of sorting equilibria in our model that will help us interpret what we have found in the data and guide our further analysis.

Suppose that positive numbers of individuals choose each sector and that

$$h^a(x, s(x)) \leq h^n(x, s(x)) \quad (12)$$

i.e., for each ability x human capital is at least as high in non-agriculture than in agriculture. Then it must be that

$$W^a = W^n \iff h^a(x, s(x)) = h^n(x, s(x)) \text{ and } \tau = 0 \quad (13)$$

that is, the rental rates of human capital are the same in both sectors if and only if the technologies that translate characteristics into human capital are the same in both sectors and the allocation of labor is efficient. In this case, all individuals are indifferent in our model and many sorting equilibria exist. We are interested in the one in which individuals with $x < x^*$ work in agriculture and individuals with $x > x^*$ work in non-agriculture. In this sorting equilibrium, the average wage is larger in non-agriculture than in agriculture:

$$\frac{W^a H^a}{L^a} < \frac{W^n H^n}{L^n} \quad (14)$$

As we will see in the next section, this case accounts for the wage gap between non-agriculture and agriculture in the U.S. In order to establish this, we will estimate W^i and show that $W^a = W^n$. The next proposition summarizes the results derived thus far:

Proposition 1. Suppose that (8)–(12) hold and that $\tau \geq 0$. If $W^a = W^n$, then:

- (i) $h^a(x, s(x)) = h^n(x, s(x))$
- (ii) $\tau = 0$

There is a sorting equilibrium in which

- (i) $H^a < H^n$
- (ii) $W^a H^a / L^a < W^n H^n / L^n$

In the other case, i.e. $W^a \neq W^n$, it is impossible in general to make statements about $h^i(x, s(x))$ and τ . However, if the technologies are the same, then there is another result that will turn out to be useful for interpreting the international results below:

Proposition 2. If (8)–(11) hold and $h^a(x, s(x)) = h^n(x, s(x)) \forall x$, then $\tau = (W^n - W^a)/W^n \in (0, 1)$.

4 Further Facts about Sectoral Wages and Human Capital in the U.S.

4.1 Mincer returns by sectors (U.S. 1980–2011)

The standard way of constructing human capital from section 2 assumes that the return of a year of schooling is the same in agriculture and non-agriculture. The model section above suggests that this may be overly restrictive. We therefore assess now whether the Mincer returns to education differ between the two sectors. To this end, we run the following regression:


$$\log(w_{ijt}) = \beta_{dt}d_t + \beta_{di}d_i + \beta_x X_{ijt} + d_i \times \beta_{si}S_{ijt} + d_i \times \beta_{ci}C_{ijt} + \varepsilon_{ijt}$$

We find that indeed some Mincer returns differ: $\beta_{sa} < \beta_{sn}$ while $\beta_{ca} \approx \beta_{cn}$. A useful way to represent this result is by showing the estimated log-wage function $\log(w_i(s))$ for each of these sectors. These are given in Figure 2. Clearly, the returns to schooling are initially steeper in non-agriculture.


There are two different ways of interpreting these differences. The “sectoral hypothesis” attributes the differences to sectoral technologies: schooling generates less human capital for workers who choose agriculture. The “selection hypothesis” attributes the differences to the workers: schooling generates the same human capital in each sector, but workers sort according to their unobserved ability. As a result, non-agriculture is more human capital-intensive and has higher Mincer returns than agriculture. To distinguish between the two hypotheses, we report in Table 5 the sectoral human capital and wages per efficiency unit that are implied by our estimates of the sectoral Mincer returns. In the first line, we calculate human capital gaps in each of our three data sources. These gaps are roughly twice as large as those computed with aggregate returns to schooling. Indeed, they are nearly as large as productivity gaps and the residual productivity gaps range from modest (in the Census) to non-existent (in the March CPS). Our preferred data set, the monthly CPS files, indicates only a 3 percent difference. Such a difference could easily be rationalized by a cost-of-living differential of the size that we have documented in previous work, Herrendorf and Schoellman (2013).

The evidence that we have just presented suggests that $W^a = W^n$. Under the assumptions of Proposition 1, our model implies that there is a sorting equilibrium in which:


$$(i) \quad h^a(x, s(x)) = h^n(x, s(x))$$



(a) U.S. Census



(b) March CPS



(c) Monthly CPS

Figure 2: Returns to Schooling Vary by Sector

(ii) $\tau = 0$

(iii) $H^a < H^n$

(iv) $W^a H^a / L^a < W^n H^n / L^n$

We should also mention that while Proposition implies that (i)–(ii) must hold in any sorting equilibrium if $W^a \approx W^n$, there are other sorting equilibria in which both (iii)–(iv) do not hold. Since (iv) is clearly born out by the data, we conclude that if $W^a = W^n$, then our model has a unique sorting equilibrium that accounts for the key features of U.S. data. Interestingly, (iv) implies in the context of our model that the allocation of labor between non-agriculture and agriculture is efficient, which is what the literature on structural transformation assumes, at least for the U.S. Implication (i) that $h^a(x, s(x)) =$

Table 5: Gaps Measured with Sector-specific Returns to Schooling (U.S. 1980–2011)

	U.S. Census	March CPS	Monthly CPS
Human Capital Gaps	0.65	0.52	0.57
Residual Wage Gaps	0.88	1.00	0.97

$h^n(x, s(x))$ is rather restrictive. In the next subsection, we provide evidence from individuals who switch sector (“switchers”) that is broadly consistent with (i).

4.2 Mincer returns for switchers (U.S. 1980–2011)

We now study the wage patterns of workers who switch sectors (“switchers”). If switchers experience large changes in the return on their education, this indicates that the differences are intrinsic to the sectoral technology. This would contradict the conclusions derived so far. If switchers do not experience changes in the return on their education, this indicates that they are intrinsic to the workers. This would support the conclusions derived so far.

Our ideal data set has a panel dimension so that we can compare the wages and wage structure before and after switching industries for the same worker. Our ideal data set also has a large sample so that we can observe a sufficient number of switchers, given the relatively small size of the agricultural labor force in the U.S. today. The only dataset that we are aware of that satisfies these requirements is the CPS. The CPS provides the sample size that we need and includes a short panel structure: households are in the CPS for four months, then out for eight months, before returning for four more months. We focus on matching the fourth month of each spell, when extra data are collected (the so-called “outgoing rotation groups”). These two observations are separated by one year. We can study the changes in wages and wage structure for workers who switch sectors in the intervening year.

Matching workers over time in the CPS is well known to be challenging. The basic problem stems from two points. First, the CPS is really a survey of addresses, not persons or households. That is, the CPS samples dwellings based on address and surveys whoever lives in that dwelling. In some cases, the family in that dwelling will differ over time. In principle the CPS carefully denotes when such of household change within a dwelling occurs. However, the second point is that the within-dwelling household and person identifiers are known to have coding errors. Fortunately, there is a well-established procedure for dealing with these issues and matching the CPS that we follow here (Madrian and Lefgren 1999,

Table 6: Wage Gaps for Matched and Unmatched Samples (U.S. 1980–2011)

	Monthly CPS	Matched CPS
Raw Wage Gap	0.58	0.53
Adjusted Wage Gap	0.55	0.49
Standard Human Capital	0.69	0.63
Sector-Specific Returns	0.97	0.96

Madrian and Lefgren 2000). The basic idea is to start by matching all persons who share the same dwelling, household, and person identifiers. We then check whether the match is valid by checking whether variable responses are logically consistent across time within a match. Here, there is a tradeoff. If one checks more variables and/or requires stricter agreement over time, then one excludes more false matches, but also excludes some valid matches where a code is misreported. We adopt a fairly strict check by requiring that age, sex, and race all agree.⁴

The matching process excludes a number of observations. Before analyzing the data, we want to make certain that the same basic facts apply also in the sample of matched workers. If so, then we can be reassured at least that workers we can match do not differ in obvious ways from the workers we exclude. We focus on two moments that are of the most interest to our analysis. First is the wage gap. Table 6 shows the raw wage gap that prevails in the monthly CPS and the matched CPS. The two should agree closely because they start with the same observations; the only difference is that the matched file contains only households that can be matched across years. Indeed the gap is close in size. The Table also shows the gap after making a series of adjustments for state and gender; for schooling, using the observed aggregate return; and for schooling, using the observed sectoral return. Throughout the estimates are close. This is reassuring because it says that the matching process yields a dataset with wage gaps similar to the baseline data.

Given our interest in the sectoral return to schooling we also check whether the basic patterns are the same for matched and unmatched data sets. We estimate the same equation allowing for sectoral returns to schooling in each dataset and plot the resulting wage as a function of schooling profiles in Figure 3. The two profiles are indeed quite similar, again reassuring us that the results we find for the matched CPS will apply to the other U.S.

⁴Agreement here means that the gender is the same; that the race report of white or non-white is the same; and that the age in the later period is between the same and two years older than the age in the earlier period. Note that since the CPS is not necessarily asked on the exact same date of the month each time, this is the strictest agreement on age that one can check.

**Table 7: Wage Changes for Stayers and Switchers
(U.S. 1980–2011)**

a→a	a→n	n→a	n→n
0.05	0.22	-0.08	0.07

data sets discussed in the previous section.




Figure 3: Sectoral Return to Schooling in Matched and Unmatched CPS Data

Having established that the matched CPS appears to be a useful dataset to study, we now turn to analyzing the experiences of switchers and non-switchers. We define four groups: those who work in agriculture both years; those who work in non-agriculture both years; those who switch from agriculture to non-agriculture; and those who switch from non-agriculture to agriculture. We focus on the determinants of the annual changes in wages. That is, we want to understand for worker i who works in sectors j and j' in periods $t - 1$ and t the determinants of $\Delta \log(w_{ijj't}) \equiv \log(w_{ijj't}) - \log(w_{ijj't-1})$. As a simple exploratory regression we try:

$$\Delta \log(w_{ijj't}) = \beta_{dj'j'} d_{jj'} + \beta_{dt} d_t + \varepsilon_{ijj't}.$$

where jj' denotes a sector pair; for example, $jj' = a, a$ denotes workers who stay in agriculture, and $jj' = n, a$ denotes workers who switch from non-agriculture to agriculture. The idea of this regression is simply to capture the mean effect of staying in the same sector versus switching, controlling for trend wage growth with d_t .

Table 7 shows the resulting estimates. Workers who remain in agriculture and non-agriculture gain roughly 5 and 7 percent higher wages. The more interesting figure is for switchers: workers who switch from agriculture to non-agriculture gain roughly 22 percent higher wages, while workers who switch in the opposite direction lose roughly 8 percent of their wages. These figures are quite small relative to the total agricultural wage gap of roughly 40 percent. We view these results as providing evidence that most of wage gaps represent selection of workers of different types. When an identical worker is moved, the wage gap moves by (on average) just 15 percent.

A common concern is selection. One can imagine that all workers have the chance to switch sectors and that there is some heterogeneity in the wage gain or loss to switching sectors. In most models, workers with larger gains (or smaller losses) will be more likely to switch. This logic suggests (for example) that the 8 percent wage loss from moving to agriculture from non-agriculture may understate the wage loss that would occur if a random sample of workers were moved. Note, however, that the same logic works in reverse: the 22 percent wage gain from moving from agriculture to non-agriculture is probably higher than the wage gain that would occur if a random sample of workers were moved. A simple bounding argument suggests that the underlying wage difference between agriculture and non-agriculture is between 8 and 22 percent.

We now ask how switching affects the market value of education. Recall that the cross-sectional estimates indicate higher returns to education in non-agriculture than in agriculture. The estimated wage-schooling relationship for the two sectors is given in Figure 4a. The relationships in the matched CPS data suggest large differences in the return to education. In Figure 4b we add the returns to education for those who switch, as observed before the switch. We can see already that switchers are different even before they switch: their wage-return on their education is between that of workers who stay in agriculture and workers who stay in non-agriculture. This evidence suggests that selection may play a role. In Figure 4c, we confirm this suspicion: the value of education after switching is still between the two sectors, and not too different from the value of education before switching.

This last fact may be better represented by Figure 5, which shows, group by group, the value of education before and after. For workers who remain in agriculture or non-agriculture throughout the value is essentially unchanged. For workers who switch sectors, the changes are modest: workers who exit agriculture earn somewhat higher returns on their schooling, while workers who switch to agriculture experience almost no change in the relevant region. These two figures indicate that much of the gap in the return to schooling is attributable to selection: workers who are going to switch are already different from




Figure 4: Sectoral Return to Schooling for Switchers and Stayers

non-switchers in the same sector. Further, when switchers switch, their returns change little.

It may be useful to provide a quantitative sense of how much the data support the sectoral differences versus selection views of why returns to schooling differ across sectors. We can perform a simple calculation that gives such a sense. The idea is to quantify the total change in the value of education that one would observe under each of these theories; to compute the actual change observed in the data; and then to compare the three values to assess how close the data is to each theory.

For the selection view, the answer is straightforward: rates of return differences reflect differences in the workers who select into the two sectors. When workers switch they take their returns with them. Hence, switchers would experience no change in the value of their




Figure 5: Wages as a Function of Schooling, Before and After Switching

education. This value is reported in column 2 of Table 8.

For the sectoral view, the answer requires some calculation. We observe in the data S_{ijt} , C_{ijt} , β_{sj} , and β_{cj} . We can compute the total value of the schooling for a worker i in sector j as $\beta_{sj}S_{ijt} + \beta_{cj}C_{ijt}$. The average value of the human capital stock for sector j is simply the average of this across all workers in j . Under the view that rates of return differences are due to sectoral factors, the counterfactual value of this education in sector j' is just $\beta_{sj'}S_{ijt} + \beta_{cj'}C_{ijt}$, with the change in the value of education defined as the difference between the two. In words, we can compute the change in the value of education for an individual by assigning them the red line instead of the blue or vice versa in Figure 4a. We compute the average change implied by this calculation for switching all agricultural workers to non-agriculture, and vice versa. The sectoral view implies large changes in the value of education that would raise or lower wages by about 50 percent.

In the fourth column of the table we give the actual wage change induced by substituting the estimated “before” returns to schooling for the estimated “after” returns to schooling for the two groups of switchers. In words, we are quantifying the effect of moving from the purple to the pink lines in Figure 6 for the average worker who actually switches. The effects are much smaller than our counterfactuals. For workers who switch from ag to non-

Table 8: Change in Value of Education

	Change Implied by:		Actual Change
	Selection View	Sectoral View	
a→n	0.00	0.54	0.23
n→a	0.00	-0.44	-0.07

Table gives the wage change due to changing value of schooling for switchers. Values are provided for two stylized theories as well as the actual change observed in the model.

ag, the effect is a 23 percent increase in wages. This is closer to the selection view. For workers who switch in the opposite direction the actual change is even smaller, just a seven percent decline in wages. We conclude that the data are much closer to the selection than the sectoral view of why returns to schooling vary by sector.

5 Cross-country Analysis

5.1 Preliminaries

In the previous sections we analyzed wage gaps in the U.S. In this section, we extend our analysis to a sample of 26 censuses from 8 countries with the relevant variables in IPUMS International. We use Censuses from IPUMS because they are nationally representative and the IPUMS team has put forth a great deal of effort in terms of harmonizing variables such as education and industry that are important for our analysis. The following country-year pairs satisfy these criteria: Brazil (1980,1991,2000); Canada (1971,1981,1991,2001); India (1983,1987,1993,1999,2004); Indonesia (1976,1995); Jamaica (1982,1991,2001); Mexico (1990,2000,2010); Panama (1970,1990,2000); Venezuela (1981,1990,2001). In addition to rich countries like Canada, our sample comprises relatively poor countries like India and Indonesia that have large workforces in the sector in which they are very unproductive, agriculture. One shortcoming of our sample is that it is fairly small and does not contain African countries. The reason for this is that IPUMS does not report sufficiently detailed census information about wages by sector for other countries than the ones in our sample.

5.2 Cross-country facts about wages and human capital

We now calculate wage gaps and assess the importance of human capital in accounting for them. We follow the same broad principles as before to the extent possible. In particular,

Table 9: Wage Gaps across Countries

	Minimum	Median	Maximum
Raw Wage Gap	0.24	0.51	0.70
Adjusted Wage Gap	0.33	0.56	0.75
Standard Human Capital	0.57	0.70	0.84
Sector-specific Returns	0.64	0.80	1.00

we use only information from individuals who are at least 10 years old and have valid information for our key variables: sex, education, employment by sector and so on. Within the set of countries in our sample we again find that there are wage gaps between non-agriculture and agriculture. Table 9 gives the raw wage gaps and the adjusted wage gaps (which again control for differences in gender and region of residence in the country). The median gap is actually quite comparable to the U.S., but there is substantial dispersion. There is also less correlation with GDP p.c. than one would expect based on previous work on agricultural productivity gaps. For example, the largest gap is for Brazil 1980, but the smallest is for Venezuela 2001. We then explore how far we can go to closing these gaps with various corrections for sectoral human capital differences. Table 9 shows that the answer is: quite far. Our preferred human capital correction closes wage gaps from 0.51 to 0.80 for the median country. Figure 6 displays the results. The x axis is the raw gaps described in Table 9. Along the y axis we show the successive effects of making the same adjustments that we implemented for the U.S.. This figure shows visually the importance of adjusting for human capital, using both years of schooling and the sectoral return to schooling. We are able to close more than half of the total wage gap in the median country. This limits the potential role of barriers/wedges. The next subsection makes that statement more precise.

5.3 Interpreting the cross-country facts

To understand the results from the previous section, we start with the observation that the literature typically assumes that except for TFP differences, technology is the same across countries. Making that assumption in our context implies that $h^a(x, s(x)) = h^n(x, s(x))$ also across countries. Under conditions (8)–(11), Proposition 2 from above then implies that $\tau = (W^n - W^a)/W^n \in (0, 1)$. In other words, viewed through the lens of our model, the gaps in residual wages give us a size of the distortion τ . Table ... reports these distortions. We can see that ...




Figure 6: Wage Gaps and Adjusted Wage Gaps

6 Discussion

6.1 Relation to the literature

Our results are related to Gollin, Lagakos and Waugh (2011), who seek to measure and account for gaps between non-agricultural and agricultural productivity. The main difference between the two papers is that Gollin et al. focus on a much larger set of countries including the poorest ones from Africa. While that makes their sample more representative than ours, the wage data by sector from population censuses that we use here do not exist for most of their countries. They therefore use Mincer returns that are common across all countries to construct measures of human capital gaps across sectors. Since we have wage data by country and sector, we use the estimated Mincer returns by sector instead. We find that the human capital gaps constructed with estimated Mincer returns are considerably larger than those constructed with the Mincer returns used by Gollin et al. (2011).

Our results are also related to Young (2013), who documented for 65 poor and middle-income countries that migration flows go in both directions: while on average one in five individuals born in rural areas moves to urban areas as an adult, one in four individuals born in urban areas moves to rural areas as an adult. Young develops a location model in which

people sort depending on observed characteristics such as schooling and on unobserved characteristics such as skills that are acquired after leaving school. The basic idea is that if observed and unobserved characteristics are positively correlated and are more important as an input in production in urban areas, then sorting implies that on average individuals with higher observed and unobserved characteristics locate in urban areas and individuals with lower observed and unobserved characteristics locate in rural areas. Since our concept of human capital is broad in the sense that it includes the stocks of observed and unobserved characteristics, this prediction of Young is consistent with our finding that most of the wage gaps between non-agriculture and agriculture are accounted for by the fact that workers in non-agriculture have considerably more human capital (broadly measured).

6.2 Implications for the Macro-Development Literature

Our work has several important implications for the macro-development literature. First, it implies that to construct human capital at the sectoral level one needs to use Mincer returns that are estimated at the sector level, instead of Mincer returns that are estimated at the aggregate level. The reason for this is that Mincer returns that are estimated at the sectoral level capture that unobserved characteristics differ across sectors, which Mincer returns that are estimated at the aggregate level miss by construction. Second, our result that in the U.S. wages per efficiency unit are equalized between non-agriculture and agriculture is a necessary condition for the efficient allocation of labor between the two sectors, which the literature on structural transformation assumes. Note, however, that it is not a sufficient condition for efficiency because it does not rule out distortions that indirectly affect the allocation of labor without affecting wages, an example being value added taxes or subsidies that affect the two sectors different. Third, our result that even in the poorer countries of our sample the potential role for distortions/barriers is relatively small limits the quantitative importance of surplus labor in agriculture ...

7 Conclusion

We have documented for eight countries ranging from rich (Canada, U.S.) to poor (India, Indonesia) that average wages are considerably higher in non-agriculture than in agriculture. We have proposed a measure of human capital at the sector level that accounts for the entire wage gap in the U.S. and most of the wage gaps elsewhere. The key feature of this measure has been that Mincer returns to schooling were estimated at the sectoral level,

instead of at the aggregate level. We have developed a multi-sector model that explains these finding if: (i) Mincer returns to schooling are equal in both sectors; (ii) more able workers sort into non-agriculture; (iii) distortions to the allocation of labor between sectors are zero in the U.S. and small elsewhere. We have argued that (iii) implies that the role for misallocation of labor between non-agriculture and agriculture is rather limited.

Our work opens several interesting avenues for future research ...

References

- Bils, Mark and Peter J. Klenow**, “Does Schooling Cause Growth?,” *American Economic Review*, December 2000, 90 (5), 1160–1183.
- Binelli, Chiara**, “How the Wage-Education Profile got More Convex: Evidence from Mexico,” 2012. mimeo, University of Southampton.
- Caselli, Francesco**, “Accounting for Cross-Country Income Differences,” in Philippe Aghion and Steven N. Durlauf, eds., *Handbook of Economic Growth*, Vol. 1A, Elsevier, 2005, chapter 9, pp. 679–741.
- Gollin, Doug, David Lagakos, and Mike Waugh**, “The Agricultural Productivity Gap in Developing Countries,” 2011. mimeo, Arizona State University.
- Hall, Robert E. and Charles I. Jones**, “Why Do Some Countries Produce So Much More Output Per Worker Than Others?,” *Quarterly Journal of Economics*, February 1999, 114 (1), 83–116.
- Herrendorf, Berthold and Todd Schoellman**, “Why is Measured Productivity So Low in Agriculture?,” 2013. mimeo, Arizona State University.
- , **Richard Rogerson, and Ákos Valentinyi**, “Growth and Structural Transformation,” in Philippe Aghion and Steven N. Durlauf, eds., *Handbook of Economic Growth*, Vol. 2, Elsevier, 2013, pp. 855–941.
- Kuznets, Simon**, “Modern Economic Growth: Findings and Reflections,” *American Economic Review*, 1973, 63, 247–258.
- Lemieux, Thomas**, “The Mincer Equation Thirty Years After *Schooling, Experience, and Earnings*,” in S. Grossbard, ed., *Jacob Mincer, A Pioneer of Modern Labor Economics*, Springer Verlag, 2006, pp. 127–145.

- Madrian, Brigitte C and Lars John Lefgren**, “An approach to longitudinally matching Current Population Survey (CPS) respondents,” *Journal of Economic and Social Measurement*, 2000, *26* (1), 31–62.
- **and Lars Lefgren**, “A note on longitudinally matching Current Population Survey (CPS) respondents,” 1999. NBER Working Paper t0247.
- Psacharopoulos, George**, “Returns to Investment in Education: A Global Update,” *World Development*, 1994, *22*, 1325–1343.
- Restuccia, Diego, Dennis T Yang, and Xiaodong Zhu**, “Agriculture and Aggregate Productivity: A Quantitative Cross-Country Analysis,” *Journal of Monetary Economics*, March 2008, *55* (2), 234–250.
- Young, Alwyn**, “Inequality, the Urban-Rural Gap, and Migration,” *The Quarterly Journal of Economics*, 2013, *128* (4), 1727–1785.

Data Appendix