Why Is Labor Productivity so Low in Agriculture?*

Berthold Herrendorf and Todd Schoellman
Arizona State University†

September 26, 2011

Abstract

Why is agriculture much less productive than the rest of the economy in many countries? A common answer to this important question in development economics is that barriers to the movement of labor or goods cause large productivity gaps. In this paper we provide evidence from US states during 1980–2009 that agriculture is much less productive in many US states. We show that the standard explanations for productivity gaps account for very little only, and that barriers play no role at all. We then provide evidence suggesting that agricultural productivity is seriously mis–measured in the United States.

*We thank Roozbeh Hosseini, Moritz Kuhn, David Lagakos, Patrick Minford, Richard Rogerson, Ákos Valentinyi, and Gustavo Ventura, as well as seminar participants at the Free University of Berlin, the 2011 Midwest Macroeconomics Meetings, the 2011 Meetings of the Society for Economic Dynamics, the 2011 Cologne Workshop on Macroeconomics, and the Universities of Toronto and Western Ontario for helpful discussions and suggestions. For financial support, Herrendorf thanks the Spanish Ministry of Education (Grant ECO2009-11165).

†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

Contents

1 Introduction

It is well known that developing countries are much less productive in agriculture than in the rest of the economy. Gollin, Lagakos and Waugh (2011), for example, review 112 developing countries and find that the median productivity gap is three while the maximum productivity gaps are around ten.[1] It is also well known that developing countries typically have the majority of their labor forces work in agriculture. Taken together, these facts imply that most people in developing countries work in the most unproductive sector. This raises two obvious questions: Why are there such large sectoral productivity gaps between non–agriculture and agriculture? Why do people not move out of unproductive agriculture into productive non–agriculture?

A common answer to these questions is that there must be barriers that distort the allocation of resources between agriculture and non–agriculture. In this spirit, Restuccia et al. (2008), for example, demonstrated that in a two–sector model with agriculture and non–agriculture barriers to the movements of labor and intermediate goods can lead to large sectoral productivity gaps even though a sizeable part of the workforce is in agriculture. What is missing from this literature to date is direct evidence that establishes the existence of barriers or of the resulting wage gaps in developing countries. The reason why there is no such direct evidence is that the available data are very limited. For example, the Penn World Table, which is the largest set of cross–country data that includes many developing countries, contains only information about final expenditure but not about employment and wages.

In this paper, we go down a different path and study productivity gaps between non–agriculture and agriculture in the fifty US states. Focusing on US states has the advantage that we can draw on rich and well documented data that is not available for developing countries. Nonetheless, some readers might think that there is not much to be learned from focusing on the United States. After all, US agriculture produces only a small percentage of total value added. Moreover, there should only be small productivity differences compared to non–agriculture, which are easily accounted for by sectoral differences in average human capital, labor shares etc. In what follows we will present evidence that challenges this view. To begin with, although US agriculture produces only a small percentage of total value added, its total size is large. For example, agricultural value added per capita in the United States is of the same order of magnitude as in Africa, although of course the share of agriculture in total value added is a lot larger in Africa than in the United States.

[1] Other recent papers which found similar results are Caselli (2005) and Restuccia, Yang and Zhu (2008).

More surprisingly, we will find that the best measurement results in large productivity gaps between non–agriculture and agriculture in many US states that are not at all accounted for by sectoral differences in average human capital, labor shares etc.

As a first pass at the evidence, we choose US data sources in 2000 that are as similar as possible to those that are available for developing countries. In particular, for sectoral value added we use the BEA's regional accounts (which form the basis for NIPA) and for sectoral employment we use workers from the Census and the Current Establishment Survey (CES). From these data sources, we calculate "conventional measures" of value added per worker in current dollars for non–agriculture and agriculture. Our conventional measures suggest that in most US states during 2000, productivity per worker in current dollar was considerably higher in non–agriculture than in agriculture. In several states the resulting productivity gap was more than a factor of three and in the maximum state it was almost a factor of ten.

Since the conventional measures of productivity gaps have several shortcomings, we move on to create the best possible measure of productivity gaps in US states. To this end, we take sectoral value added from the United States Department of Agriculture (USDA) and sectoral hours worked in primary and secondary jobs from the CPS. We calculate value added per hour worked in non–agriculture and agriculture at the state level for the thirty–year period 1980–2009. To have enough observations, we group the data into ten–year bins over which we take averages. Averaging over ten–year bins takes care of the concern that our results are driven by a few bad harvests in agriculture. We also exclude five states with very small agricultural sectors. We find that while the conventional measurement inflates the productivity gaps somewhat, fairly sizeable productivity gaps remain with the best measurement. Moreover, we find that these gaps do not at all decline over time. To give some concrete numbers, we find that the median productivity gap is 1.6, the productivity gap at the 90th percentile is 2.8, and the maximum productivity gap is 5.7. So even the best available measurement leads to the surprising conclusion that there are large and sustained productivity gaps for US states during 1980–2009.

These findings raise the obvious question about what may account for the productivity gaps between non–agriculture and agriculture in US states. To answer that question, we first develop a decomposition of the productivity gaps into the four most common explanations: sectoral differences in human capital, differences in the cost of living, barriers to the movement of labor, and sectoral differences in labor shares. We then use various data sources to measure as carefully as we can each of these explanations. We emphasize that this is possible because the United States have exceptionally detailed data such as the

American Time Use Survey, which we use extensively.

We find that there are quantitatively large sectoral differences in human capital at the state level, and that they account for a sizeable part of the productivity gaps. In contrast, the sectoral differences in the cost of living and the size of barriers are quantitatively small and account for hardly anything. The fact that that barriers turn out to be unimportant is consistent with the common view that US economy is undistorted. We also find that there are large differences in the labor shares within agriculture; agriculture in coastal state and in states with less agriculture tends to be much more labor intensive whereas agriculture in the Great Plains tends to be much more capital intensive. To our surprise, however, these differences in the agricultural labor shares do not help at all to account for the gaps in productivity. In fact, taking them into account undoes almost everything that the previous three explanations achieved: taken together the four explanations account only for a negligible part of the productivity gaps. The reason why sectoral differences in labor shares have this effect is that to account for productivity gaps between non–agriculture and agriculture, agriculture would need to have a larger labor share than non–agriculture (so that a larger share of the smaller value added goes to labor). In contrast, the data show that agriculture tends to be more capital intensive than non–agriculture in most states, and particularly capital intensive in states with large productivity gaps.

Since the standard explanations account only for a negligible part of the observed productivity gaps between non–agriculture and agriculture, we are left with the uncomforting possibility that either the productivity gaps or at least one of the standard explanations are badly measured. Although our initial prior was that this is very unlikely given the alleged quality of US data, our results force us to take the possibility of mis–measurement seriously, and so we explore it to the extent possible by comparing our preferred measures with alternative measures. While we find no evidence that the potential explanations have serious measurement problems, this exercise leads us to conclude that both value added and hours worked in agriculture are very badly measured in US states. For example, we show that calculating agricultural value added at the state level from different source leads to quantitatively sizeable discrepancies that can be as large as a factor ??. We take this to imply that mis–measurement of agricultural productivity is likely to be problem in developing countries, at least to the extent that they have lower quality than US data. We conclude from this that improving the measurement of agricultural productivity in developing countries is of first–order importance in order to establish robust facts that can guide future research on sectoral productivity gaps.

Our paper joins a growing literature about sectoral productivity differences across countries, which aims to identify the sectors that make poor countries unproductive. While some of this literature explores more than two sectors (e.g., Herrendorf and Valentinyi (2010)), most papers focus on the apparent gap between agricultural and non–agricultural productivity. Our paper is most closely related to Gollin et al. (2011) who also try to measure and account for the gap between agricultural and non–agricultural productivity given the existing data. The main difference between the two papers is that Gollin et al. (2011) use surveys from developing countries. While this provides more direct evidence regarding productivity gaps, it comes at the costs that these surveys tend to be considerably less detailed and reliable than the US data that we use. In any case, from the perspective of this paper, it is telling that Gollin et al. (2011) also find it challenging to account for the observed productivity gaps. Our paper is also linked to a line of research that seeks to explain where sectoral productivity gaps come from. Existing theories point to the scale or risk of farming (Adamopoulous and Restuccia 2010, Donovan 2011); to barriers between agriculture and non–agriculture (Restuccia et al. 2008, Herrendorf and Teixeira 2009); to differences in factor endowments (Caselli 2005); to selection in the workers in the two sectors (Lagakos and Waugh 2010); and to home production (Gollin, Parente and Rogerson 2004).

Hsieh-Klenow on firm level variation.

The remainder of the paper proceeds as follows. Section 2 measures sectoral differences in productivity first from conventional data sources that would be available across countries and then from the best available data that are only available for US states. Section 3 provides some basic theory on how to decompose sectoral productivity gaps and then connects the theory to the data. Section 4 discusses measurement error as an explanation for our findings. Section 5 concludes. An appendix contains a detailed description of our data sources.

2 Measurement of Productivity Gaps

In this section, we document that productivity in non–agriculture is higher than in agriculture in most US states. Productivity is defined as value added in current dollars per worker or per hour. Since the ratio of productivity in non–agriculture to agriculture is unit free, we can compare it across states and time. If this ratio is not equal to one, then we say that there is a sectoral *productivity gap*. We will first measure sectoral productivity gaps using the conventional data sources that are of the sort that is typically available for developing countries. Afterwards, we will measure sectoral productivity gaps by using the best available data sources for US states.

2.1 Conventional Measurement of Productivity Gaps

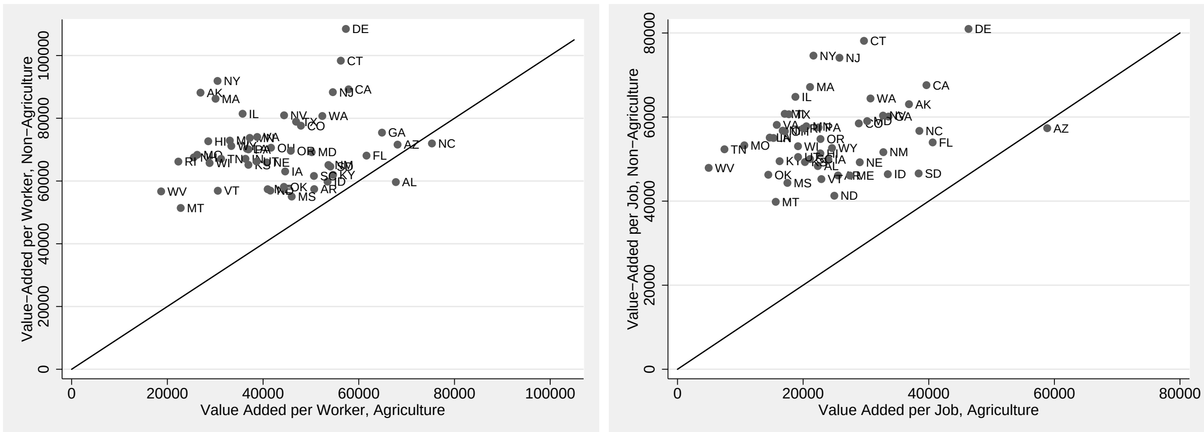

(a) Census Employment

(b) CES Employment

Figure 1: Sectoral Value Added per Worker in 2000

We define the agriculture sector as the farm industries crop and animal production. Our definition of agriculture is consistent with that of the Food and Agricultural Organization (FAO). It does not include fishing, horticulture, hunting, and trapping, and it includes forestry only to the extent to which it is done on farms. We define the non–agriculture sector as all industries other than the farm industries and the military. We exclude the military from the non–agricultural sector because we do not have employment data by state for it.

We use data sources for US states that are as similar as possible to the standard data sources that are available for developing countries. For agricultural value added in current prices, the standard sources are the United Nations National Accounts and the World Development Indicators (WDI), which are both based on national accounts. The closest comparable numbers for US states are the value added numbers by industry and state from the BEA's regional accounts. For employment, the standard data source is the International Labor Organization (ILO), which is based on population censuses or on employment surveys. The closest comparable numbers for US states are the employment–by–industry numbers from the Population Census or from the Current Establishment Survey (CES), which underlies NIPA. We obtain the Census numbers from the public–use version made available through Ruggles, Alexander, Genadek, Goeken, Schroeder and Sobek (2010) and the CES numbers from the BEA's regional accounts.[2]

[2]In 1997, the industry classification in the BEA's regional accounts switched from SIC to NAICS. We

Table 1: Conventional Measurement of Productivity Gaps in 2000

	BEA/Census	BEA/CES
Median	1.7	2.4
90$^{\text{th}}$ Percentile	2.8	3.6
Maximum	3.3	9.7

To take a first look at the data, we initially focus on the year 2000, which is the most recent census year available. Figure 1 shows the productivity gaps for 2000. Panel (a) refers to Census employment and panel (b) refers to CES employment. Agricultural value added per worker is depicted on the x–axis and non–agricultural value added per worker is depicted on the y–axis. We can see that for almost all states productivity is higher in non–agriculture than in agriculture. Upon closer inspection, we can also see that the resulting productivity gaps are sizeable and change considerably depending on whether we use Census or CES employment numbers. Table 1 reports some summary statistics. The table confirms the previous impression: the productivity gap in the mean state equals 1.7 and 2.4 depending of which measure of agricultural employment we use and with CES labor all summary statistics of productivity gaps are larger than with Census labor.

In sum, we find that there are sizeable productivity gaps between non–agriculture and agriculture for the year 2000. We also find that these gaps are sensitive to the way in which sectoral employment is measured. These results suggest that more careful measurement is required. In the next subsections, we will use the best data sources that are available for the United States. We emphasize that these data sources are of higher quality than what is typically available for developing countries.

2.2 Improved Measurement of Employment

For the improved measurement, we use employment data from the *Current Population Survey (CPS)*, which are generally thought to be of the highest quality available for the United States. Comparing the CPS data with those from the Census and the CES, we will show that Census and CES data have serious shortcomings that affect the measurement of sectoral employment at the state level. We will also show that it makes a quantitative

merge the two series so that all estimates up to 1997 are based on SIC categorizations and all estimates from 1998 onward are based on NAICS categorizations. The Census changes its industry coding schemes, which also differ from those of the BEA. The Appendix gives the crosswalk between Census industries and our non–agricultural and agricultural sectors.

Table 2: Workers (in millions) by Sector in 2000

	Census	CPS	CES
Agriculture	1.6	2.4	3.1
Non–Agriculture	127.8	136.9	161.5

Table notes: Economy–wide employment in millions. CPS numbers measured as full time equivalents.

difference whether we calculate sectoral productivity as value added per worker or as value added per hour worked.

The CPS is a rotating panel survey administered by the BLS in each month of the year to a sample of households. It contains information on the identities and hours worked of the first two jobs of one quarter of respondents in each month from January 1994 onward. The CPS also contains surveys on multiple job holdings during May of 1979, 1980, 1985, 1989, and 1991, giving us a relatively good coverage before 1994.

Table 2 compares employment by sector at the national level from the Census, the CPS and the CES in 2000. Recall that the Census and the CES numbers are those that we have used for our conventional measurement above. The CPS numbers are for full–time equivalent employment, i.e., 48 weeks and 40 hours/week. The table shows that CPS employment lies between Census and CES employment in both sectors. To the extent that the CPS is deemed to be the most reliable source of employment data, this implies that Census numbers underestimate and CES numbers overestimate sectoral employment. The table also shows that the relative discrepancies between the Census and CES numbers are considerably larger in agriculture than in non–agriculture. This explains why in the previous subsection the measured productivity gaps came out larger when we used CES employment.

There are two main reasons for the differences between the employment numbers from the Census and the CES. The first reason has to do with how the two data sources deal with secondary jobs. The Census records only primary jobs, implying that it underestimates employment in both sectors. In contrast, the CES counts all jobs including secondary ones, but at the state level it does not adjust secondary jobs for full–time equivalent employment, implying that it overestimate sectoral employment.[3] While in principle this could be equally important in both sectors – or more important in either one of them – it turns out that the relative discrepancies between the two data sets are larger in agriculture than in non–

[3] While the BEA makes this adjustment at the national level, it not clear how successful it is at this; see Cociuba, Prescott and Ueberfeldt (2009) and Bureau of Economic Analysis (2010, chapter IX).

agriculture, because there is a sizeable number of workers with first jobs in non–agriculture and second jobs in agriculture. The second reasons for the differences between the employment numbers from the Census and the CES is that the Census is taken during the month of March. Again this is more important in agriculture than non–agriculture, because, as Figure 2 shows, employment in agriculture is highly seasonal and March is a month with below average activity in agriculture. The figure also shows that non–agricultural employment is hardly seasonal.[4]

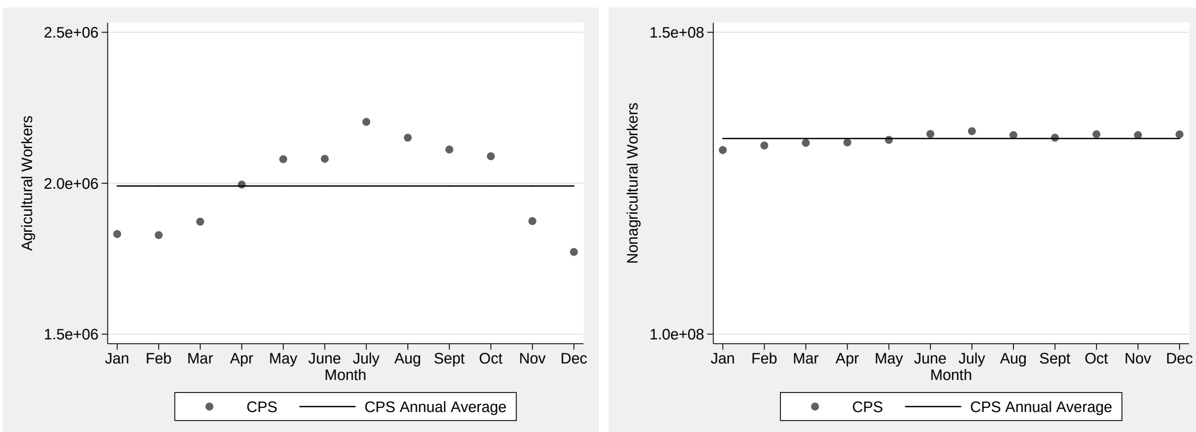

(a) Agriculture

(b) Non–Agriculture

Figure 2: Employment during 2000

An additional potential issue with the standard data sources is that they tend to report employment by sector but not hours worked by sector. If hours per worker are roughly the same in both sectors, then using employment instead of hours worked does not affect the calculations of productivity gaps. The information of hours worked by sector from the CPS allows us to assess whether this condition holds for the United States. We find that it does not hold, and that workers in agriculture tend to work more hours. For example, restricting attention to workers with only one job, the average farmer works as much as 40.1 hours whereas the average non–farmer works just 37.1 hours. Not taking the sectoral difference in hours worked into account leads to an underestimation of sectoral productivity gaps. Table 3 compares the summary statistics for value added per worker based on employment versus hours worked. The first two columns are repeated from Table 1 above to make comparisons easier. The third column reports the sectoral productivity gaps calculated

[4]Note that if we used hours, we would also find that seasonality is much more important in agriculture than in non–agriculture. Note too that even usual hours turn out to be highly cyclical in agriculture, which is largely due to the fact that the respondents in different months are different workers.

Table 3: Productivity Gaps in 2000 – Improved Measurement

	BEA/Census	BEA/CES	BEA/CPS	USDA/CPS
Median	1.7	2.4	2.1	1.8
90^{th} Percentile	2.8	3.6	3.5	3.2
Maximum	3.3	9.7	5.0	4.5

as the ratio of BEA value added (as before) and CPS hours worked (new). We can see that using hours worked increases the productivity gaps compared to using Census labor and reduces them compared to using CES labor. So we conclude that sizeable sectoral productivity gaps result also when we measure employment as hours worked instead of workers.

2.3 Improved Measurement of Value Added

The numbers for farm value added that we have used above were from the BEA's regional accounts. A potential issue with them is that for some reason the BEA chooses to report the value added produced by *farmers*, i.e., persons who operate a farm, instead of the value added produced on *farms*.[5] To appreciate the significance of using farmers instead of farms, an example may be helpful. Consider the payments that are received by farm contractors or the rental payments that are received by land owners who are not farmers. Clearly, these are factor payments that are produced on farms, and so they belong to the value added produced on farms. However, the BEA does not report them as part of farm value added because they do not lead to income of farmers.

To see how much the difference between the value added produced by farmers and the value added produced on farms matters for measuring productivity gaps, we construct a new measure of value added that includes all factor payments generated on farms irrespective of whether they accrue to farmers or non–farmers. To this end, we go to the State–Net–Value–Added Accounts provided by the USDA, which is the original data source of the BEA. Since these accounts include a detailed breakdown of the receipts, expenses, and factor payments in the farm sector at the state level, they have sufficient information to make the required adjustments. The Appendix describes in detail how we use this information to calculate the value added produced on farms.

Figure 3 plots the comparison between value added from the BEA and the USDA at the state level in 2000. All value added numbers in the figure are per CPS hour worked; the

[5]Note that the USDA uses the same concept of farm value added as the BEA.

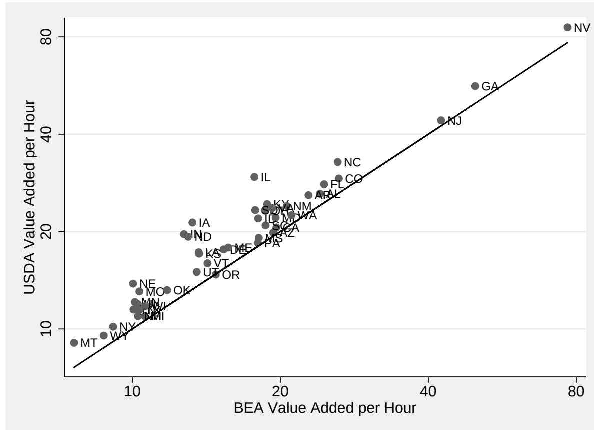

Figure 3: Agricultural Value Added by State in 2000 from Different Data Sources

numbers based on the BEA are on the x–axis and the numbers based on the USDA are on the y–axis. Expectedly, the figure shows that all observations are above the 45 degree line, that is, value added per hour worked based on the USDA is larger in all states than that based on the BEA. Perhaps somewhat unexpected, the figure also shows that for states with low agricultural productivity the difference can be large quantitatively. This suggests that using the BEA numbers for value added will bias downwards productivity in agriculture, and thereby will inflate the measured productivity gaps between non–agriculture and agriculture in a typical state. Columns three and four of Table 3 above report the summary statistics for value added per hour worked based on BEA and USDA value added. We can see that, as expected, the productivity gap shrink when we replace BEA value added with USDA value added. However, even for USDA value added, the productivity gaps remain sizeable, and each of them is larger than the corresponding summary statistic in the first column for BEA value added and Census labor.

In sum, the previous discussion has established two key points. First, the best available measure of productivity gaps, which uses USDA value added and CPS hours worked, reports that there are sizeable productivity gaps at the level of US states in the year 2000. Second, these gaps are considerably smaller than if one uses BEA value added and CES employment. Since BEA value added and CES employment are similar to the standard data sources that are available for developing countries, this suggests that at least part of the large productivity gaps that people find for developing countries are likely to be the result of flaws in the data sources that are available.

2.4 Improved Measurement of Productivity Gaps

We are now ready to provide the most careful measurement of sectoral productivity gaps in the United States for the longest time period for which there is data. As suggested by the discussion in the previous two subsections, this is going to be based on value added from the USDA and hours worked from the CPS. The longest period for which there is data in the CPS is 1978–2009. To make sure that we have enough observations in the CPS, we group the data in three non–overlapping ten–year bins: 1980–1989, 1990–1999, and 2000–2009, which we casually refer to as the 1980s, the 1990s, and 2000s. All numbers we report are the ten–year averages of the sectoral productivity gaps by state within the respective bins. Even with ten–year bins, the CPS has only a few farmers for some states that have small agricultural sectors. To make sure that our results are not driven by these states, we require that for all states in our sample the CPS have complete hours information for at least 90 agricultural workers in each of the three decades. This criterion leads us to exclude Alaska, Connecticut, Massachusetts, Rhode Island, and West Virginia.

Panel (a) of Figure 4 gives the results in the form a histogram. We can see that with the best possible measurement and thirty years worth of data the stylized fact of Section 2.1 survives; in most states and years, there are considerable productivity gaps, that is, productivity in non–agriculture is considerably higher than in agriculture. Panel (b) of Figure 4 shows that the productivity gaps do not decline over time. This, and the fact that we average over ten–year bins, addresses the concern that our results above reflected a bad harvest in several states during the year 2000. Panel (c) of Figure 4 shows the time series for the five states with the largest productivity gaps. Note that these are the five states with the largest productivity gaps after we excluded Alaska, Connecticut, Massachusetts, Rhode Island, and West Virginia. If we had left them in, then Alaska and West Virginia would be among the states with the largest productivity gaps.

Table 4 gives the summary statistics for the productivity gaps in US states that we obtain with improved measurement: the median gap is 1.6, the gap at the 90$^{\text{th}}$ percentile is 2.8, and the maximum gap is 5.7. While these gaps are somewhat smaller than the ones we found for the 2000, we need to keep in mind that we have excluded five states from the sample. Nonetheless, the summary statistics are surprisingly close to those in developing countries. For example, Gollin et al. (2011) document for a set of 112 developing countries that the median productivity gap is 3 and the 95$^{\text{th}}$ percentile is 8.8.

In what follows, we take the gaps of Table 4 as the baseline and ask what accounts for them. The next section provides some basic theory about the reasons for productivity gaps, and then it connects the theory to the US data.

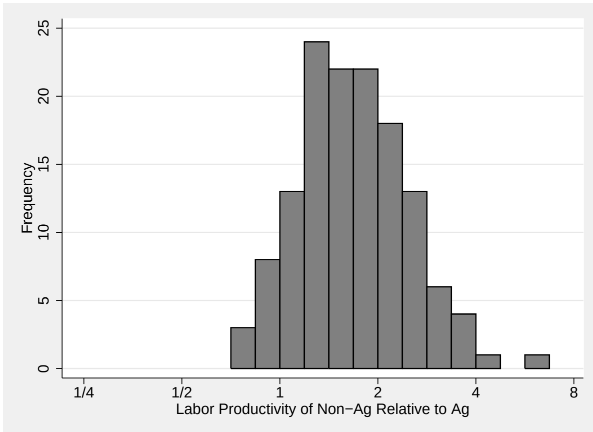

(a) Histogram

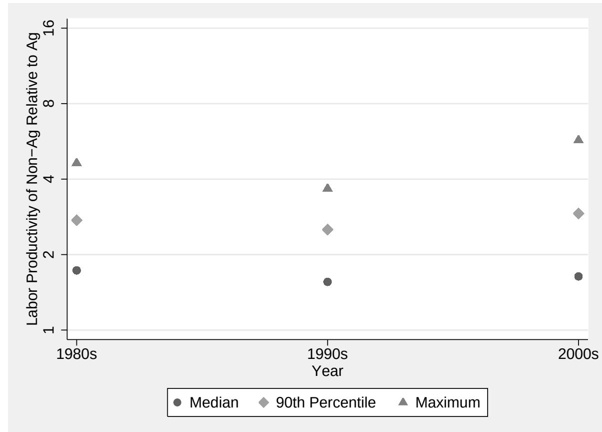

(b) Panel

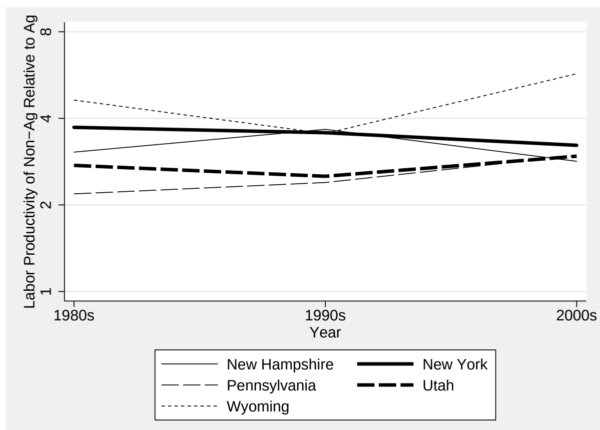

(c) Time Series for States with Largest Gaps

Figure 4: Improved Measurement of Labor Productivity Gaps 1980–2009

3 Decomposition of Productivity Gaps

3.1 Theory of Productivity Gaps

We start by developing a simple model of productivity gaps that spells out the key causes for them. Our model is a static partial equilibrium model that takes several variables as exogenously given. We choose not to write down a fully articulated general equilibrium model, because doing so would increase complexity without adding insights that are important for our project here.

Consider an economy with one period and S states. Each state has two separate locations. In one location agricultural production takes place and in the other location non–

Table 4: Productivity Gaps 1980–2009 – Improved Measurement

	USDA/CPS
Median	1.6
90^{th} Percentile	2.8
Maximum	5.7

Table notes: US results are for 45 states, excluding five states with small samples: Alaska, Connecticut, Massachusetts, Rhode Island, and West Virginia.

agricultural production takes place. We index states by $s \in S$ and locations by $j \in \{a, n\}$, with a and n standing for agriculture and non–agriculture.

Productivity per hour worked in location sj is defined as

$$\frac{Y_{is}}{l_{sj} N_{sj}},$$

where Y is value added in current dollars, l is average hours worked per worker, and N is the number of workers. Denoting the dollar wage per hour by W, the following identity links wages and productivity per hour worked:

$$W_{sj} = LS_{sj} \frac{Y_{sj}}{l_{sj} N_{sj}}, \tag{1}$$

where

$$LS_{sj} \equiv \frac{W_{sj} l_{sj} N_{sj}}{Y_{sj}}$$

is the labor share in value added. Equation (1) simply says that the wage in a sector equals the share of value added that goes to labor times the value per hour worked. Note that equation (1) is an identity that holds irrespective of the market structure or of the properties of the production function. As a result, it does not explicitly feature distortions that work at the margin, which will be convenient when we connect it to the data. Note also that one can choose whether or not subsidies are included in the measure of value added that enters equation (1), as long as the choice is the same for the labor share and productivity per hour worked.

To obtain an expression involving the productivity gap, we divide equation (1) for locations sn and sa by each other and rearrange:

$$PG_s \times \left(\frac{W_{sn}}{W_{sa}}\right)^{-1} \times \frac{LS_{sn}}{LS_{sa}} = 1, \tag{2}$$

where

$$PG_s \equiv \frac{Y_{sn}/(l_{sn}N_{sn})}{Y_{sa}/(l_{sa}N_{sa})}.$$

We will call the left–hand side of (2) the adjusted productivity gap. Intuitively, the adjusted productivity gap equals one because productivity gaps must be reflected in equal gaps in wage payments.

Two remarks about the decomposition (2) are at order. First, we emphasize that there is no reason to view $PG_s = 1$ as the benchmark case. Instead, (2) shows that even if the dollar wages per hour are equalized across sectors, sectoral differences in labor shares cause sectoral productivity gaps. The reason, of course, is that given a common dollar wage for both sectors, the sector with the lower labor share needs to have higher productivity in order to be able to pay that common wage. Second, the way in which we have written the decomposition (2) deliberately imposes a strict measure of success that requires us to account for productivity gaps on a state–by–state basis. Specifically, (2) starts with a productivity gap for a given state and then cumulatively applies the adjustments for that state. The merits of this approach can best be explained with an example. Suppose that we had only two states with productivity gaps of 1 and 2 and with cumulative adjustments of 1/2 and 1. If we make the adjustments according to (2), then the adjusted productivity gaps turn out to be 1/2 and 2, and so the maximum productivity gap remains unchanged. If instead we wrote the decomposition as

$$PG_s = \frac{W_{sn}}{W_{sa}} \times \left(\frac{LS_{sn}}{LS_{sa}}\right)^{-1},$$

then we would conclude that the maximum adjustment of 2 equals the maximum productivity gap. While that might appear to be success, it is not because the maximum productivity gap comes from the first state whereas the maximum adjustment comes from the second state.

The next step is to explain how sectoral wage gaps arise. To this end, we need to model the household side. We assume that in each location, there are many households which derive utility from consuming a final good but do not value leisure. Households are endowed with productive hours and stocks of physical capital, land, and human capital. These endowments vary across households.

The timing within the period is as follows. At the beginning of the period, each household chooses whether to stay in its location or to move to the other location in its state. Afterwards, it supplies all its productive hours in the chosen location and its capital and land in economy–wide rental markets. Then, the household earns wages and rental payments and purchases the final good at the dollar price $\mathcal{P}_{sj}$ that prevails in its location. We will call $\mathcal{P}_{sj}$ the cost of living in location sj.

The natural starting point is to study this economy without barriers to the movement of labor. In this benchmark case, households are indifferent between the two locations of the state in which they live. Since they derive utility only from the consumption of the final good, this implies that real wages per efficiency unit are equalized across the locations of each state:

$$\frac{W_{sn}}{\mathcal{P}_{sn} h_{sn}} = \frac{W_{sa}}{\mathcal{P}_{sa} h_{sa}},$$

where, as before, W_{sj} denotes the dollar wage per hour worked and h_{sj} denotes average human capital. Rearranging this equation gives:

$$\frac{W_{sn}}{W_{sa}} \times \left(\frac{h_{sn}}{h_{sa}} \times \frac{\mathcal{P}_{sn}}{\mathcal{P}_{sa}} \right)^{-1} = 1. \tag{3}$$

In words, if there are no barriers to the movement of labor, then a gap in the dollar wage per hour reflects differences in human capital and the cost of living. This implies that locations with higher human capital or higher cost of living have higher hourly wages.

The alternative to the benchmark case is that there are barriers to the movement of labor. To capture this, we assume that moving between the two locations of state s requires a dollar cost of $B_s > 0$. In this case, there are two types of equilibria: either households in each locations prefer strictly the location in which they are or households in one location prefer strictly where they are while households in the other location are indifferent. Since we see people moving between agriculture and non–agriculture, we focus on the second type of equilibrium. The equilibrium gap in the dollar wage per hour then satisfies

$$\frac{W_{sn}}{W_{sa}} \times \left(\frac{h_{sn}}{h_{sa}} \times \frac{\mathcal{P}_{sn}}{\mathcal{P}_{sa}} \times \mathcal{B}_s \right)^{-1} = 1, \tag{4}$$

where

$$\mathcal{B}_s \equiv \begin{cases} 1 + B_s & \text{if } \frac{W_{sn} h_{sa} \mathcal{P}_{sa}}{W_{sa} h_{sn} \mathcal{P}_{sn}} > 1 \\ \dfrac{1}{1 + B_s} & \text{if } \frac{W_{sn} h_{sa} \mathcal{P}_{sa}}{W_{sa} h_{sn} \mathcal{P}_{sn}} < 1 \end{cases}.$$

In words, the observed gap in the dollar wages per hour reflects barriers to the movement of labor, in addition to differences in human capital and the cost of living.

Note that sectoral wage differences may be due to reasons other than sectoral differences in human capital, the cost of living, and barriers. For example, wages in agriculture may be higher than in non–agriculture to compensate for the fact that work in agriculture is more physically demanding. The structure that we have imposed on the decomposition implies that these kinds of sectoral differences will be picked up by barriers.

Combining equations (2) and (4), we arrive at the decomposition of sectoral productivity gaps that we will connect to the data:

$$PG_s \times \left(\frac{h_{sn}}{h_{sa}} \times \frac{\mathcal{P}_{sn}}{\mathcal{P}_{sa}} \times \mathcal{B}_s\right)^{-1} \times \frac{LS_{sn}}{LS_{sa}} = 1. \qquad (5)$$

This decomposition makes explicit that sectoral productivity gaps may be due to sectoral differences in four variables: human capital, the cost of living, barriers, and labor shares.

3.2 Empirical Results

We are now ready to measure the importance of each of the four variables for sectoral productivity gaps. As before, we focus on the three ten–years bins during the period 1980–2009 and the 45 US states in our restricted sample. We start by measuring the sectoral differences in human capital, the cost of living, barriers, and labor shares.

3.2.1 Measuring gaps in human capital

To construct a measure of human capital in each sector and state, we use CPS data on wages, age, education, and gender. We follow standard practice and focus on the restricted sample of workers who have one job, work at least 30 hours per week, are employed for wages and salaries, and have between zero and fifty years of potential experience (i.e., age minus years of schooling minus 6).[6] We run separate regressions for agriculture and non–agriculture at the national level, with log wages being the dependent variable and dummies for age, school degree, and gender being the dependent variables. We use the regression results to predict a wage per worker for each ten–year bin. We average these predicted wages over the workers in each sector, state, and ten–year bin. Our measure of human capital in a sector and state during the ten–year interval is the average predicted wage from this procedure minus the intercept from the wage regression.

Several remarks are at order. First, we run the log–wage regressions only for the restricted sample of workers with good wage information, but use the predicted wages for all workers. This is common practice at least since Katz and Murphy (1992). Second, we do

[6]Note that although these sample restrictions reduce the number of agricultural workers from our earlier sample, we still have at least 39 observations remaining in each state during each decade.

not include the intercepts in our measure of human capital by sector, because they capture the average unexplained wage level that may be due to reasons other than human capital (e.g., unionization in parts of nonagriculture). Third, we allow the rates of return to differ by sector. In contrast, the development literature typically uses the same economy–wide rate of return for both sectors; see for example Hall and Jones (1999) or Caselli (2005). Figure 5 provides the justification for allowing rates of return to differ by sector; the rates of return to schooling in agriculture are lower than in nonagriculture and the age–wage profile is flatter in agriculture. This is consistent with the view that workers self–select into the two sectors and it suggests that less human capital accumulation on the job takes place in the agriculture than in non–agriculture. Our measure of human capital captures these features of the data.

Following the procedure explained above, we find that there are large differences between human capital in non–agriculture and agriculture at the state level: the human–capital gaps of the states at the 10^{th} percentile, the median, and the 90^{th} percentile of the human–capital gap distribution are 1.43, 1.81, and 1.92. This means that in the median state, an average worker in nonagriculture has almost twice the human capital of an average worker in agriculture.

Figure 5: Wage Profiles by Sector, 2000 CPS

(a) School–Wage Profiles (b) Potential Experience–Wage Profiles

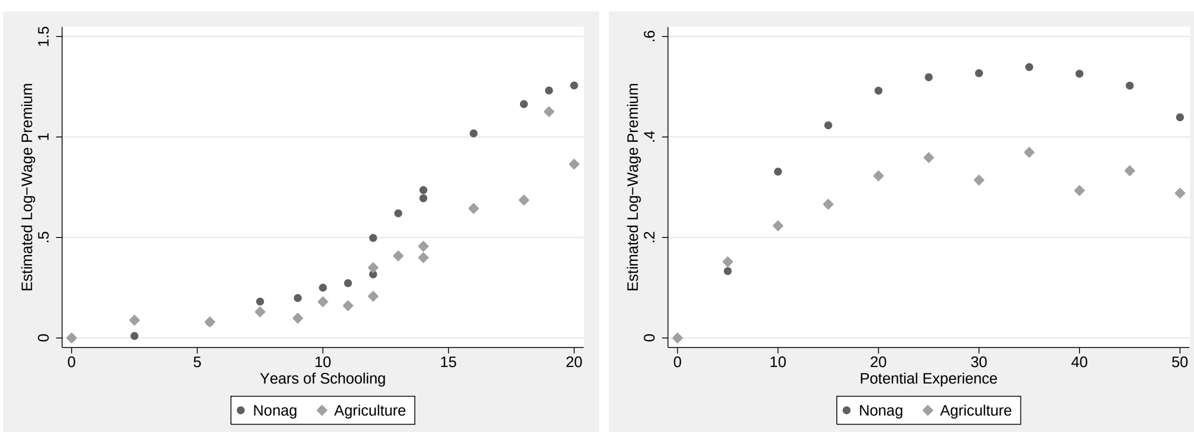

3.2.2 Measuring the gaps in the cost of living

We continue with measuring the sectoral gap in the cost of living between non–agriculture and agriculture at the state level. The cost of living has two dimensions: locations differ in commute times to work and in the cost of living. We will measure both dimensions.

We take commute time into account by calculating the effective number of hours which is defined as the sum of hours commuted and hours worked. We calculate commute time in each sector as the average commute time per day times the commutes per worker per week times the workers per week. The average commute time per day is from the Census.[7] The commutes per worker per week are from ATUS for 2003–2009; we calculate the average and use it for all years. The workers per week are from the CPS and are available for all years.

We calculate the cost of living in a sector as the average price level that the workers of this sector face. To calculate the price level for a given worker, we need to know the sector in which he works and the price level which he faces at his residence. The information about the sector and the residence comes from the Population Censuses. As for commute time, we use the numbers from the 1980 Census for 1980–1989 etc. Given our prior discussion, one might think that it would be preferable to use CPS data instead of Census for this, but before 1986 the geographic detail in the CPS is too limited and the sample size is too small to produce reliable estimates for smaller metropolitan ares. Since this is not an issue with the Population Censuses, we use them instead. The information about the price level in the residence of a given worker is harder to come by with than one might imagine, because the CPI covers only major metropolitan areas. Moreover, the CPI is normalized to 100 in the base year, which makes level comparisons impossible. Fortunately, recent research by Aten (2006), Aten and D'Souza (2008), and Aten (2008) has constructed consumer price levels for 363 metropolitan areas as well as for the rural area of each state for the year 2006. We combine this information with the residence information to obtain a price level for each worker. We then average across the price levels of the workers of each sector and state so as to obtain the average price level in that sector and state.[8]

We find that across the entire United States the gap in the price levels between metropolitan and non–metropolitan areas is 1.44. While this price gap is large, a sizeable part of it comes from price variation across states, whereas the price gap in the median state of the price gap distribution is only 1.29. Even this number is not the relevant one for our purposes, however, because some agricultural workers live in metropolitan areas (typically smaller ones) and many rural residents do not work in agriculture. When we factor this in we find that gaps in the price levels between non–agricultural and agricultural workers are sizeable only in the few states that have large, expensive urban areas. For example, the gap is as large as 1.34 in New York and 1.23 in Illinois and Virginia. In contrast, the gap

[7] We use the numbers from the 1980 Census for 1980–1989, from the 1990 Census for 1990–1990, and from the 2000 Census for 2000–2009.

[8] Appendix C contains a more detailed description of how we calculate the price levels.

in the price levels of the median state of the price gap distribution amounts merely to 1.08.

3.2.3 Measuring barriers

This will change after benefits are taken into account.

Barriers to the movement of labor are not directly observable. To measure barriers we use the fact that, as we have shown above, they lead to sectoral wage gaps. Since dollar wages per hour in non–agriculture and agriculture are directly observable at the state level in the CPS, it is fairly straightforward to measure sectoral wage gaps. Using them together with our measures of sectoral gaps in human capital and in the cost of living, we can then infer the size of barriers from equation (4). Figure 6 shows the histogram of the real wage gaps per efficiency unit that result after making the adjustments for differences in human capital and the cost of living. We can see that the real wage gaps per efficiency unit are somewhat to the left of 1, yet close to 1. The implied barriers for the 10$^{\text{th}}$ percentile, the median, and the 90$^{\text{th}}$ percentile of the distribution of barriers are 0.7, 0.9, and 1.0, respectively. The fact that $\mathcal{B}$ comes out somewhat smaller than one in the median state indicates small barriers for moving from non–agriculture to agriculture in the median state. To interpret this result, recall that in our framework barriers pick up all sectoral differences other than human capital and the cost of living. Since there do not seem to be physical barriers for moving to agriculture, our results are reflect the fact that agricultural work is more physically demanding and dangerous, which needs to be compensated by a higher wage in agriculture.

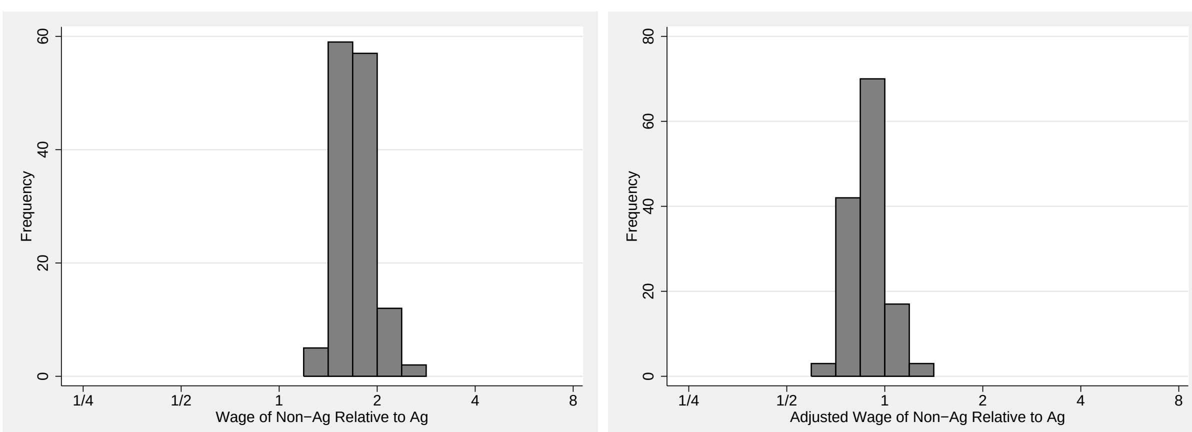

Figure 6: Wage Gaps

3.2.4 Measuring gaps in labor shares

We start with the labor share in non–agriculture. Since agriculture value added amounts only to a very small part of aggregate value added, we set the labor share in non–agriculture equal to the aggregate labor share. Since the BEA does not break down value added at the level of the different states, it is not obvious how to estimate the labor shares at the state level. **Plot the wage by state in non–agriculture against the 0.67 times the productivity in non–agriculture. Hopefully get the 45 degree line.** We therefore use the nation–wide aggregate labor share and set $LS_{sa} = 0.67$ for all states.

It is important to realize that in agriculture the nation–wide aggregate labor share would not be a good proxy. The first reason is that the nation–wide labor share in agriculture is considerably smaller than the nation–wide aggregate labor share [Herrendorf, Rogerson and Valentinyi (2011)]. Moreover, different regions of the United States specialize in different types of agriculture (think of fruits and vegetables versus cereals and grains), and these different types of agriculture are likely to have different labor intensities, which may have important implications for our results. We therefore calculate the labor shares in agriculture at the state level. Note that this is feasible because there is much richer data for agriculture at the state level than for non–agriculture.

Our method of calculating agricultural labor shares uses the observed factor inputs and market rental rates and then imputes the factor payments to land, physical capital, and labor in agriculture for each state. To impute the payments to agricultural land, we draw on the information in Turner, Tamura, Schoellman and Mulholland (2011), who report for each state the acreages of three types of agricultural land (i.e., irrigated cropland, other cropland, and pastureland) and the corresponding dollar rental rates. To impute the payments to physical capital in agriculture, we again draw on the information in Turner et al. (2011), who report also the dollar values of physical capital in agriculture in each state. We obtain the payments to physical capital by multiplying these numbers with a standard economy–wide dollar rate of return. To impute the payments to labor, we use the information on the total cost of hired farm labor in each state from the Census of Agriculture. This leaves the payments to self–employed and unpaid workers for the labor that they provide. We could use the self–reported information on wages of self–employed workers, but we have the usual concern that these wages are unreliable. Instead, we therefore impute the wage payments for both self–employed and unpaid workers from the information contained in the CPS. Appendix D contains a detailed description of the steps involved in different imputations.

We find that there is substantial cross–state variability in the labor shares of agriculture. In particular, the imputed labor share ranges from 0.18 in Iowa to 0.70 in New Hampshire,

Table 5: **Adjusted Sectoral Productivity Gaps**

	PG	h	$\mathcal{P}$	$\mathcal{B}$	LS
Median	1.6	1.0	0.9	1.0	1.5
90$^{\text{th}}$ Percentile	2.8	1.8	1.4	1.6	2.2
Maximum	5.7	4.2	4.1	3.2	4.7

Table notes: Results are for the forty-five states for 1980–2009. Results are cumulative, that is, last column reflects all adjustments.

and the labor share in the median state is **Median labor share in agriculture**. In general, coastal states and states with less agriculture tend to have higher labor shares whereas states in the Great Plains tend to have lower labor shares. This is consistent with the casual observation that states like California, Florida and Hawaii specialize in the labor intensive production of coffee, fruits or vegetables, whereas the states in the Great Plains specialize in the capital intensive production of grains and animals.

The question for this project is whether the considerable variation in the agricultural labor shares that we find for US states will help to explain why there are such large sectoral productivity gaps. For that to happen, the labor shares in agriculture would need to be higher than in non–agriculture in states with large productivity gaps. The next subsection assesses whether this is the case.

3.2.5 Decomposing productivity gaps

Table 5 reports the results of the decomposition. From left to right, the results are cumulative. The table shows that taken together the five explanations account only for a negligible part of the productivity gaps. Put differently, we find that almost the entire productivity gaps are unexplained residuals.

Looking at the explanations separately, we can see that differences in sectoral human capital do in fact account for a sizeable part of the productivity gaps, whereas differences in the cost of living and barriers only play a negligible quantitative role. Perhaps surprisingly, labor shares widen the productivity gaps and undo a sizeable part of the effect of the other explanations. This is the case because agriculture is more capital intensive than the rest of the economy. The fact that agriculture has a lower labor share than non–agriculture implies that, everything else the same, productivity in agriculture should be higher than in non–agriculture, not lower. Although that is indeed the case in some states, it turns out that these are not the states with the large productivity gaps.

We should stress that after the adjustments the distribution hardly changes,

suggesting that measurement error is a serious possibility.

4 Measurement Error

In this section, we will ask why there are such large unexplained residuals of the sectoral productivity gaps. An obvious starting point is that we have left out an explanation that will account for them. Although this is possible, of course, we emphasize that we have explored all standard explanations that are being considered in the literature. Therefore, the more likely explanation for our results is that there is serious mis–measurement of the sectoral productivity gaps and/or the explaining factors. In what follows, we will provide evidence that supports the view that mis–measurement is an important factor behind our results.

4.1 Measurement of gaps in wages and in labor shares

We start by discussing how reliable our measurement of the wage gaps is. Our measures of wages are based on micro data from the CPS which are generally deemed very reliable. **We need more discussion here.**

We continue by discussing how reliable our measurement of the labor shares is. To provide additional evidence that our estimates are in fact reasonable, we estimate the agricultural labor shares at the state level also according to the different method that was proposed by Gollin (2002).[9] The basic idea of Gollin was to calculate the labor share only for the part of value added that can be assigned unambiguously to capital or labor. Essentially, this boils down to subtracting proprietors income from value added (because it contains payments to both capital and labor) and then splitting the residual among capital and labor. Obviously this method does not work so well if the share of proprietors income in value added is large. Since this is the case for a few US states, we have not used Gollin's method as our preferred method, but we still think that it can serve to provide additional evidence that the range of agriculture labor shares across US states is indeed large. Applying Gollin's method to the 45 states of our sample, we find two comforting results: (i) The labor shares from our method and his method are highly correlated, with a correlation coefficient of 0.89; this suggests that there is substantial agreement between the two methods about which states have more or less labor intensive agriculture; (ii) the estimates from Gollin's method display even *more* variation than those from our method;

[9] Appendix D contains a detailed description of what we do and which data we use.

the labor shares resulting from Gollin's method actually range from a low of 0.16 in Iowa to a high of 0.86 in Hawaii (compared to a low of 0.18 in Iowa and a high of 0.70 in New Hampshire from our method).

We could repeat the exercise by focusing on employed workers only, that is, measure gaps in CPS wages of employed workers and use gaps in labor shares of employed workers. That would address the concern that everything that is related to proprietors is badly measured.

Given that we have concluded that the gaps in wages and in labor share appear to be well measured, we are left with the possibility that the productivity gaps themselves are mis–measured. The common view is that statistical agencies in the United States do a fairly good job at measuring value added per hour in non–agriculture. In contrast, it is unclear how well they do at measuring value added per hour in agriculture. In what follows, we will provide evidence that indeed there are serious measurement problems in agriculture, at least at the state level.

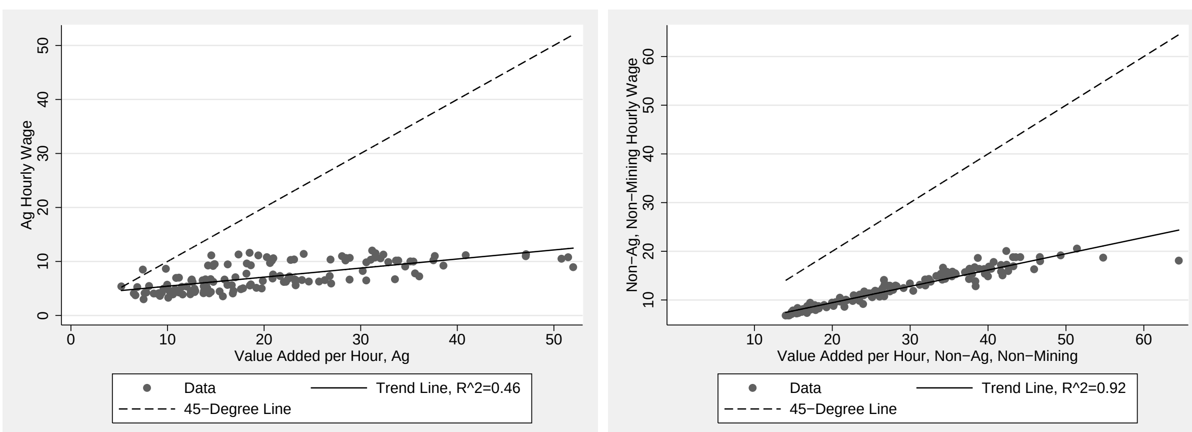

Figure 7: Gaps in Productivity and Wages during 1980–2009

Multiply value added with labor shares in figure above

4.2 Measurement of value added

We start by comparing wages per hour with value added per hour separately for non–agriculture and agriculture at the state level. The first–order condition (1) predicts that given nation–wide labor shares in the two sectors, wages per hour should be proportional to value added plus subsidies per hour. Figures 7 plot value added on the x axis against wages on the y axis for both sectors. A data point is for a state and a year during 1980–2009.

Since prices changed during this period, we do not compare the levels but focus on whether the observations are proportional to each other. We can see from the figures that while the observations line up closely in non–agriculture, they do not line up at all in agriculture, particularly for low values of value added. What is more, in several states wages per hours exhaust or even exceed value added per hour, which is conceptually impossible. This suggest that either wages per hour or value added per hour are mis–measured in agriculture. Since wages are observed in the market, the mis–measurement is likely to be with respect to value added per hour. We now examine the possibility of mis–measurement separately for value added and hours worked.

Two pieces of evidence suggest that there is mis–measurement of agricultural value added. The first one emerges when we compare agricultural value added from different data sources. Above we used the USDA because it is available every year. An alternative data source is the Census of Agriculture, which is available every five years. We compare value added constructed from these two data sources for 2002, which is the closest year to our usual comparison year 2000 for which a Census of agriculture was taken. Figure 8 plots the result. We find that the value added numbers from the Census of Agriculture range from 47–136% of those from the USDA, with the 10$^{\text{th}}$ and the 90$^{\text{th}}$ percentiles being at 74% and 118% of those from the USDA. This suggests that agricultural value added is measured with a lot of noise.

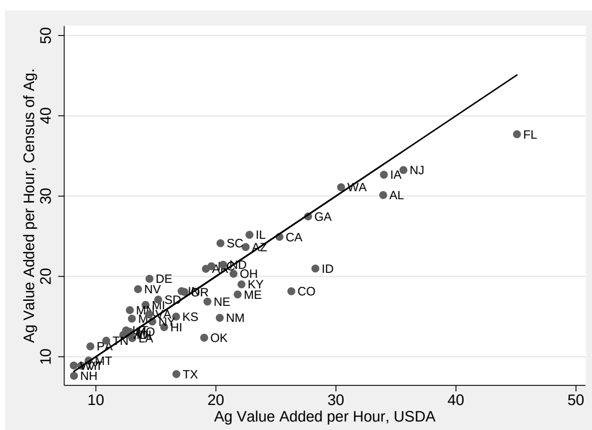

Figure 8: Value Added from USDA and Census of Agriculture

The second piece of evidence suggesting mis–measurement of agricultural value added results when we compare farm income from the Population Census and the USDA. Farm income from the Population Census is the self–reported income of farm operators, which does not include the income earned by owners who are not operators (such as corporate shareholders). Farm income from the Population Census is available for 1980 and 1990.[10] Farm income from the USDA equals the difference between farm revenues minus costs and factor payments, and so it includes both the incomes of operators and of owners who are not operators. Therefore, farm income from the Population Census should be lower than farm income from the USDA. Figure 9 plots the farm income from the two data source against each other. We can see that for several states with relatively low farm incomes, the farm income from the Population Census is considerably larger than that from the USDA, although conceptually this cannot be the case. This is consistent with the view that the USDA underreports value added in agriculture.

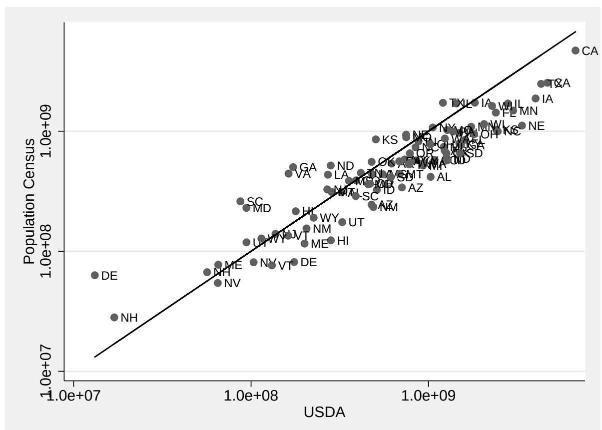

Figure 9: Operator Income from USDA and Population Census

We could compare imputed value added with measured value added also.

4.3 Measurement of hours

We also provide evidence that hours worked in agriculture are mis–measured because workers do not correctly remember how many hours they work ("recall bias"). To establish that recall bias in agriculture is an issue for the CPS, we compare the CPS numbers for hours worked with those from the American Time Use Survey (ATUS). ATUS is deemed very reliable, because its data are collected from selected CPS participants who keep time–use diaries of their activities by minute for a 24–hour period. ATUS data are available for the period 2003–2009, during which we also have CPS data. Figure 10 plots average hours per worker in agriculture and non–agriculture from the CPS against those from ATUS during 2003–2009. The disaggregation is at the level of regions because the ATUS sample has too

[10] In 2000 the Population Census stopped asking the relevant question.

few people working in farming to obtain representative estimates at the level of individual states (the numbers of farmers per region are listed in parenthesis). The figures show that ATUS hours tend to be lower than CPS hours in both sectors. This is expected because ATUS has a stricter definition of what counts as working time.[11] The figures also show two additional facts: according to both measures, agricultural hours by region vary more than non–agricultural hours; according to ATUS agricultural hours vary more than according to the CPS. It turns out that the discrepancies come mostly from the hours that farm operators report, whereas the hours that wage workers report are much more similar in both data sets. These findings cast doubt on the reliability of the hours worked that the CPS reports for agriculture.

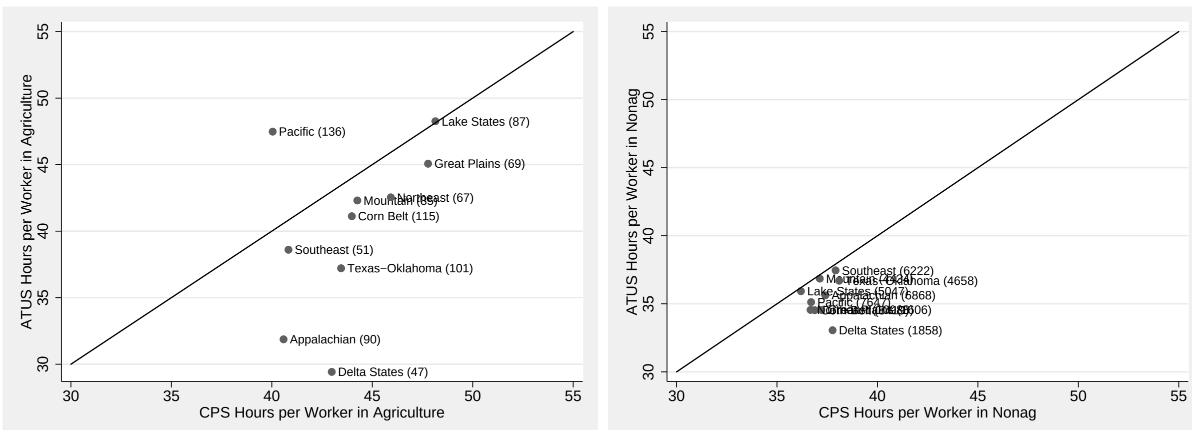

Figure 10: Hour Worked in Agriculture and Non–Agriculture from the CPS and ATUS during 2003–2009

We continue with the possible mis–measurement of the explaining factors. Decomposition 4 shows that dollar wages per hour account for the effects of sectoral differences in human capital and the cost of living and for the effects of barriers to the movement of labor. This is important in the context of mis–measurement because our measures of human capital, cost of living, and barriers could be very noisy, but dollar wages per hours are based on market prices that are known to workers. Therefore, the reported dollar wages per hour are likely to be fairly reliable. In particular, there is no obvious reason why workers in agriculture should systematically under–report and workers in non–agriculture should systematically over–report their wages, which would be needed to obtain a larger wage gap and to reduce the unexplained residual.

Although dollar wages per hour are unlikely to be mis–measured, there are still two

[11] For example, ATUS does not count business lunches or dinners as working time, but the CPS does.

potential sources of measurement error. The first one comes from the fact that the wages in the CPS do not include non–wage compensation. To the extent that non–wage compensation differs between agriculture and non–agriculture, the true wage gap may be larger than the wage gap that we have measured from the CPS. To rule out this possibility, we go to NIPA and compare wages/salaries with total compensation by industry. We find that total compensation is about 21% higher than wages/salaries in non–agriculture for the median state and 14% higher in agriculture. The second potential source of measurement error is that illegal workers could be under–represented in the CPS. To the extent that illegal workers are more important in agriculture than in non–agriculture and earn lower wages than legal workers, the true wage gap may again be larger than the wage gap that we have measured from the CPS. While that is a possibility, we note that a sizeable number of workers in the CPS report lower than the minimum wage. Moreover, the states with the five largest productivity gaps in our sample are Alaska, Connecticut, Massachusetts, Rhode Island, and West Virginia. These are not the states where migrant farm labor from is most prominent.

5 Conclusion

We studied labor productivity in agriculture and non–agriculture in U.S. states during 1980–2009. We found that there are large productivity gaps of similar size as observed in many developing countries. We also found that taken together, the standard explanations account only for a negligible part of the gaps, and barriers to the movement of labor play no role at all. We provided evidence suggesting that there are serious measurement problems in agriculture even in the United States, suggesting that even the best measures of productivity in agriculture do not give reliable results.

Our findings have several important implications for the development literature. First, we provided an example in which sizeable gaps in human capital lead to sizeable gaps in productivity while barriers play no quantitatively important role. Moreover, the fact that we find sizeable gaps in human capital across non–agriculture and agriculture for US states is remarkable because in the United States average human capital in agriculture is fairly high (most people have high school degrees). This suggest that sectoral gaps in human capital should even be larger in developing countries most of the workers in agriculture are illiterate and most of people with at least a high school degree are in non–agriculture. Second, we provided an example in which the best measurement results in large productivity gaps across non–agriculture and agriculture that are not accounted for by gaps in wages and labor shares. We argued that this points to measurement error of agricultural productivity. If even in the United States agricultural productivity cannot be measured precisely, then chances are that measurement problems are even worse for development countries. This casts some doubt on the validity of the stylized fact from the development literature that there are very large productivity gaps between non–agriculture and agriculture in developing countries.

References

Abraham, Katharine G., Sarah M. Flood, Matthew Sobek, and Betsy Thorn, "American Time Use Survey Extract System: Version 1.0 [Machine-readable database]," Maryland Population Center, University of Maryland, College Park, Maryland and Minnesota Population Center, University of Minnesota, Minneapolis, Minnesota 2008. AVailable online at `http://www.atusdata.org`.

Adamopoulous, Tasso and Diego Restuccia, "The Size Distribution of Farms and International Productivity Differences," 2010. Manuscript, University of Toronto.

Aten, Bettina H., "Interarea Price Levels: an Experimental Methodology," *Monthly Labor Review*, 2006, *129* (9), 47–61.

——, "Estimates of State and Metropolitan Price Parities for Consumption Goods and Services in the United States, 2005," October 2008. Manuscript, Bureau of Labor Statistics.

—— **and Roger J. D'Souza**, "Regional Price Parities: Comparing Price Level Differences Across Geographic Areas," *Survey of Current Business*, 2008, *88* (11), 64–74.

Bureau of Economic Analysis, "Local Area Personal Income and Employment Methodology," Technical Report April 2010, chapter IX.

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.

Cociuba, Simona, Edward C. Prescott, and Alexander Ueberfeldt, "U.S. Hours and Productivity Behavior Using CPS Hours Worked Data: 1947-III to 2009-III," November 2009. Manuscript, Arizona State University.

Donovan, Kevin, "Risk, Farm Ownership, and International Productivity Differences," 2011. Manuscript, Arizona State University.

Gollin, Doug, David Lagakos, and Mike Waugh, "The Agricultural Productivity Gap in Developing Countries," 2011. Manuscript, Arizona State University.

—— **, Stephen L. Parente, and Richard Rogerson**, "Farm Work, Home Work and International Productivity Differences," *Review of Economic Dynamics*, 2004, *7*, 827–850.

Gollin, Douglas, "Getting Income Shares Right," *Journal of Political Economy*, April 2002, *110* (2), 458–474.

Hall, Robert E. and Charlies 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 Akos Valentinyi, "Which Sectors Make Poor Countries So Unproductive?," *forthcoming: Journal of the European Economic Association*, 2010.

—— **and Arilton Teixeira**, "Barriers to Entry and Development," *Forthcoming: International Economic Review*, 2009.

—— **, Richard Rogerson, and Ákos Valentinyi**, "Growth and Structural Transformation," Chapter prepared for the Handbook of Economic Growth 2011.

Katz, Lawrence F. and Kevin M. Murphy, "Changes in Relative Wages, 1963-1987: Supply and Demand Factors," *The Quarterly Journal of Economics*, February 1992, *107* (1), 35–78.

Lagakos, David and Mike Waugh, "Specialization, Agriculture, and Cross–Country Productivity Differences," 2010. Manuscript, Arizona State University.

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.

Ruggles, Steven, J. Trent Alexander, Katie Genadek, Ronald Goeken, Matthew B. Schroeder, and Matthew Sobek, "Integrated Public Use Microdata Series: Version 5.0," Minneapolis: University of Minnesota 2010.

Turner, Chad, Robert Tamura, Todd Schoellman, and Sean E. Mulholland, "Estimating Physical Capital and Land for States and Sectors of the United States, 1850-2000," February 2011. Manuscript, Arizona State University.

A Data Appendix: Measuring Sectoral Value Added and Factor Shares

In this section, we explain in detail how we construct our measures of agricultural value added at the state level. Given these measures, it is straightforward to obtain measures of nonagricultural value added by subtracting agricultural value added from state GDP as reported by the BEA.

A.1 BEA

The BEA numbers for sectoral value added are taken straight from the BEA's regional economic accounts. In particular, for value added in agriculture, we use item 10010 (value added of farms) for years with the SIC classification and item 4 (value added of crop and animal production) for years with the NAICS classification; for GDP at the state level we use item 0 in the SIC and item 1 in the NAICS, minus value added in the military (item 112000 in the SIC and item 80 in the NAICS); for value added of mining we use item 30000 in the SIC and item 5 in the NAICS.

A.2 USDA

To construct a measure of agricultural value added from USDA data, we use the USDA's value–added spreadsheets at the state level.[12] We construct income produced on farms as follows:

- The value of crop production is farm income. The USDA reports values for eight types of crops, as well as total values for home consumption and inventory adjustment.
- The value of livestock production is farm income. The USDA reports values for four types of livestock, as well as values for home consumption and inventory adjustment.[13]
- Revenues produced from miscellaneous farm activities may or may not be counted as farm income. Considering each in turn:
  - The value of machine hire and customwork is farm income, because it includes payments for providing services closely related to the farm. Examples are planting, plowing, spraying, or harvesting for others.

[12] The spreadsheets are available at `http://www.ers.usda.gov/Data/FarmIncome/Zip_filesXls.htm`.
[13] Note that NAICS uses animal production and USDA uses livestock production for the same industry.

- The value of forest products sold from the farm is farm income. Ideally we would exclude this revenue from agriculture and include it in forestry. However, some of the expenses of farms and some of the labor on farms are devoted to generating this revenue. Since we cannot subtract off the relevant expenses and labor, we include the revenue as farm income.
- Other income is farm income, because it is closely related to farm operations. Examples include animal boarding, breeding fees, and energy generated on the farm.
- The gross imputed rental value of farm dwellings is not farm income, because it is not closely related to farm operations.

We construct expenses for intermediate inputs used by farms as follows:

- Farm–origin expenses are farm expenses. The USDA reports feed purchased, livestock and poultry purchased, and seed purchased in this category.
- Manufactured inputs are farm expenses. The USDA includes fertilizers and lime, pesticides, petroleum fuel and oils, and electricity in this category.
- Other purchased inputs may be farm expenses or factor payments. If they are factor payments, then they are not counted as farm expenses. In particular:
  - Repair and maintenance of capital items are farm expenses, in line with usual NIPA procedures.
  - Expenses for machine hire and custom work are farm expenses.
  - Marketing, storage, and transportation expenses are farm expenses.
  - Contract labor is a factor payment to contractors or crews that provide labor to farms, and so is not counted under farm expenses.
  - Miscellaneous expenses are farm expenses. Examples include the costs of animal health care and insurance.

In addition to a "raw" value added measure, we also construct a value added measure that includes subsidies. To do so we take value added and add the line "direct government payments". Note that the government also provides indirect support for farmers through price supports and similar programs; these effects are already counted in value added.

B Data Appendix: Constructing Sectoral Labor Inputs

This section provides the details of how we construct sectoral labor inputs. It also explains how we define agricultural and nonagricultural workers, hours worked, wages, and so on.

B.1 Current Employment Survey

The BEA reports employment figures for each state and industry. These figures are drawn from the Current Employment Survey. The BEA uses the same industry classification scheme for both value added and employment. Hence, we categorize workers into agriculture or non–agriculture in the same manner as with value added; see Appendix A.1.

B.2 Population Census

To calculate labor input from the U.S. Population Census, we use the public–use census data made available through IPUMS (Ruggles et al. 2010). We use the 5 percent sample for 1980, 1990, and 2000, which is the largest publicly available sample in these years. We impose as little sample selection as possible. We require that workers be in the labor force and employed (empstat = 1). The Census asks about the employment status of those aged 16 and older, so we restrict our sample to this age group. We exclude workers with invalid or missing industry or occupation codes.

To assign workers to agriculture and nonagriculture, we need to construct a crosswalk from workers' self–reports of industry/occupation to our agriculture/nonagriculture classification. One potential complication is that the U.S. Census uses a different coding scheme for occupation and industry in every year. Fortunately, however, these schemes are reasonably detailed. Throughout we work with the original classification schemes from the Censuses, rather than the re–codings performed by IPUMS ("occ" and "ind" rather than say "occ1950" and "ind1950").

Our primary form of classification draws on which industry workers report. Using the reported industries, we construct the labor force in the farm sector, i.e., animal and crop production industries. Table 6 gives the full crosswalk. It lists for each year all industries (and corresponding industry codes) that we use. We construct non–agriculture as the residual, that is, all workers with valid industry reports that do not work in an agricultural industry. We use a similar approach when we identify the workers who are not in agriculture or mining.[14]

As a robustness check on our results, we also experiment with two alternative methods of defining workers in the agriculture sector. In the first, we still use workers' reported industry codes, but we take a somewhat broader view of which industries should be counted as agriculture. In particular, we include industries identified as agricultural services or support activities for agriculture. The codes are reported in the third column of Table 6. In the second alternative, we use workers' reported occupations to identify agricultural and non–agricultural workers. That is, we identify workers who report being farmers, ranchers, farm managers, or farm laborers rather than those who report being in the animal and crop production industries. Again, the Census includes a measure of occupation that varies with each Census and generally becomes more detailed over time. Table 7 gives the occupation titles and corresponding codes that we associate with the agriculture sector for each Census year. In general both of these methods lead to similar or higher employment figures in agriculture, which in turn implies lower value added per worker in agriculture and larger sectoral productivity gaps. Results are available upon request.

After dividing the population into agricultural and non–agricultural workers, we calculate employment and hours worked by state and sector. For all calculations, we restrict the samples to individuals with valid responses. We use the reported state of residence (statefip) and weight all variables with individual weights (perwt). We compute sectoral employment as the number of workers in each sector.

B.3 Current Population Survey

The Current Population Survey (CPS) administered by the BLS is our principal data source for the number of workers and hours worked and the wages and earnings received.[15] We restrict our attention to those workers in the CPS who have a job and valid occupation and industry codes. We use information about age, education, employment, gender, and state of residence from the CPS Basic Monthly Data, which we take from the NBER's CPS data repository. We also use information about hours worked in primary and secondary jobs by industry, which is available in the May supplements of the 1980, 1989, and 1991 CPS (again taken from the NBER's CPS data repository) and in the Outgoing Rotation Groups during 1994–2009. Lastly, we use information about wages and earnings from the NBER Matched Outgoing Rotation Groups (MORG) during 1980–1993 and the Outgoing

[14]The industry codes for mining are fairly straightforward: codes 40–50 in 1980 and 1990, and 37–49 in 2000.

[15]Data are available at `http://www.nber.org/data/cps_basic.html`.

Table 6: Coding Census Industries to the Agriculture Sector

Year	Narrow	Broad = Narrow, Plus:
	Agricultural Production, Crops (010)	Agricultural Services, except Horticultural (020)
1980	Agricultural Production, Livestock (011)	
	Agricultural Production, Crops (010)	Agricultural Services, n.e.c. (030)
1990	Agricultural Production, Livestock (011)	
	Crop Production (017)	Support Activities for Agriculture and Forestry (029)
2000	Animal Production (018)	

Table notes: Numbers in parentheses correspond to codes for the variable ind in the IPUMS data.

Table 7: Coding Census Occupations to Agriculture Sector

Year	Agriculture	
	Farmers, Except Horticultural (473)	Managers, Farms, Except Horticultural (475)
1980	Supervisors, Farm Workers (477)	Farm Workers (479)
	Supervisors, Related Agricultural Occupations (485)	Graders and Sorters, Agricultural Products (488)
	Farmers, Except Horticultural (473)	Managers, Farms, Except Horticultural (475)
1990	Supervisors, Farm Workers (477)	Farm Workers (479)
	Farm Laborers and Farm Foremen, Allocated (480)	Supervisors, Related Agricultural Occupations (485)
	Graders and Sorters, Agricultural Products (488)	
	Farm, Ranch, and Other Agricultural Managers (20)	Farmers and Ranchers (21)
2000	Supervisors/Managers of Farming, Fishing, and Forestry Workers (600)	Graders and Sorters, Agricultural Products (604)
	Miscellaneous Agricultural Workers, Including Animal Breeders (605)	

Table notes: Numbers in parentheses correspond to codes for the variable “occ” in the IPUMS data.

Rotation Groups during 1994–2009.

To measure the number of workers, we use the CPS Basic Monthly Data . The CPS uses the coding schemes from the Population Censuses for both first and second jobs. The codes for agriculture are 017 during 1980–1982, 010–011 during 1983–2002, and 0170–0180 during 2003–2009. We assign each worker to the sector of his primary job, which the CPS defines as the job with most hours worked. We then count the number of primary workers in each sector, state, month, and year and weight them by the given weight (pwsswgt).

To measure the number of hours, we need to assign to each sector the hours worked in primary and secondary jobs. We therefore consider only workers with valid information on hours in their primary and secondary jobs. Before 1994 these are the workers in the special May supplements. Since 1994, these are the workers in the Outgoing Rotation Groups, which include approximately one fourth of the sample in the forth and the eight month. We multiply the resulting numbers by four to correct for sample loss. We measure hours worked as the actual hours worked in the primary and possibly secondary job during the previous week (pehract1 and pehract2). We multiply the reported weekly hours worked by 4.33 to generate the typical monthly hours worked.

The resulting data have two shortcomings. First, before 1994, hours worked are only available for the May supplements of the 1980, 1989, and 1991 CPS. Second, the reported hours worked in agriculture are sometimes based on fairly small samples. We address these issues in the following way.

1. For each sector, state, and month, we calculate the number of workers with a primary job in that sector.

2. For the months available, we calculate the total hours worked in each sector and state (where total hours are the sum of the hours worked in primary and secondary jobs).

3. For the months available, we calculate the ratio between total hours worked and the number of primary workers in each sector and state. We run a regression of this ratio on a time trend and on months effects so as to obtain the predicted total hours per primary worker in each sector, state, and month.

4. We multiply the predicted hours per primary worker by the actual number of primary workers in each sector, state, and month and aggregate the result to years. We use these predicted hours in each sector, state, and year.[16]

[16]Note that this means that we also use the predicted hours when we actually have hours in the CPS. The reason is that the sample sizes for agriculture tend to be on the small side, and so we want to smooth the data.

We measure wages per hour and perform wage regressions using the wage data from the NBER Matched Outgoing Rotation Group (MORG) during 1980–1993 and from the CPS Basic Monthly Data during 1994–2009. We include only workers in outgoing rotation groups when wage data are collected. We again multiply the weights by 4 to correct for the loss in sample size. As is standard, we run the wage regressions for a selected sample of workers who meet the following criteria:

- They work only a single job (the reason for this restriction is that the CPS does not allocate wages and salaries between different jobs).
- They works for wages and salaries (the reason for this restriction is that the reported wages of self–employed or unpaid workers are considered unreliable).
- They have a valid hourly wage. Workers can report a valid wage in two ways in the CPS: they can report a positive hourly wage or a positive weekly wage and provide a positive estimate of their usual weekly hours worked; in the latter case we compute the hourly was as weekly wage/hours per week.
- They are strongly attached to the labor market, which is measured as usually working at least 30 hours per week.
- They have between 0 and 50 years of potential experience, which is defined as age minus years of schooling minus 6.

We use the following controls in our wage regression: the state of residence, gender, potential experience, and education. We transform potential experience into 5–year bins (0–4 years, 5–9 years, and so on) and run wage regressions with dummies. Education data is straightforward except that there is a shift in the coding scheme for education in the middle of this period. Until 1991 the scheme counted years in school (such as four years of college), while from 1992 onward it measured degree attainment (such as bachelor's degree). We run wage regressions with dummies for years before 1991 and with dummies for degree from 1992 onwards.

B.4 American Time Use Survey

Our last data source for hours worked is the American Time Use Survey, which is available online from IPUMS as ATUS–X (Abraham, Flood, Sobek and Thorn 2008). ATUS is a time–use study funded by the BLS that has been fielded by the U.S. Census Bureau since 2003. It is organized as a follow–up to the CPS. Some members of each outgoing rotation group (those in their eighth month in the sample) are invited to continue into ATUS. If they choose to do so, they are interviewed again 2–5 months after their final month in the CPS. Some of the demographic information available in the CPS is collected again in ATUS. The key new data about time use are collected in a time diary, in which respondents are asked to record how they use their time during one 24–hour period from 4am to 4am. They record their time usage in minutes, selecting from over 400 detailed activities.

Our extract includes data from the years 2003–2009. For the most part our sample selection is the same as in the CPS. We require that workers be employed, whether at work or absent from it at the time of the interview (empstat = 1 or 2). We exclude workers with invalid or missing industry or occupation codes for their first job.

A complication arises from the fact that ATUS does not collect the industry and occupation of the second job. More specifically, at the time of the final CPS interview, they collect information on whether or not the respondent works multiple jobs, as well as the industry, occupation, and hours worked in up to two jobs. At the time of the ATUS interview they collect information on whether or not the respondent works multiple jobs and the industry and occupation of their first job, but not of a possible second job. Hence, we do not know from ATUS records whether a second job is in agriculture or nonagriculture. To address this problem, we keep only workers who either work one job at the time of the ATUS interview, or work multiple jobs at the times of the CPS and the ATUS interviews and list the same industry and occupation for their first job. In this case we assume that the industry and occupation for the second job did not change either. The crosswalk from industry and occupation codes in the ATUS to farm/nonfarm status is the same as in the CPS for 2003–2009 (see previous section).

In addition to the standard demographic variables, we use two pre–defined time–use variables. The first is "work, main job", which includes the minutes in the diary day spent on activities at work, checking e–mail, making phone calls to/from clients, putting in overtime, or receiving on–the–job training. This category explicitly excludes "work–related" activities, such as socializing, eating, drinking, or playing golf with coworkers or clients. It also excludes travel for most workers, unless the travel is an essential part of business, such as for taxi–drivers or traveling salespersons. The second pre–defined time–use variable we use is "work, other job(s)", which includes the same information as "work, main job" for all other jobs.

We aggregate this information into the time spent working in agricultural and nonagricultural jobs in the average week. We translate minutes per day into hours. We then aggregate this information to find the hours worked per agricultural worker, exactly as we did in the CPS. We take into account the sampling probabilities by using the provided weights (wt06).

C Cost of Living

To estimate the annual hours spent commuting for agricultural and non–agricultural workers in each state, we use that the Population Censuses asked respondents how much time they spent on a typical morning commute in the prior week. As expected, these figures are lower in agriculture than in non–agriculture, although commute times have been rising in both sectors. We multiply the morning commute by two to estimate the typical daily commute time. We then estimate from ATUS the typical number of days per week that agricultural and non–agricultural workers do any work at all. This figure is estimated at the national level since ATUS sample sizes are too small to provide state level estimates. When combined with Census information, we arrive at an estimate of the typical weekly commute time. We multiply the weekly estimate by 52 weeks in a year to arrive at the typical annual hours spent commuting for agricultural and non–agricultural workers in each state.

To construct the cost of living in agriculture and non–agriculture in each state, we use price parities that (Aten 2006, Aten and D'Souza 2008, Aten 2008) estimated for the year 2006 for 363 metropolitan areas and for the rural area of each state. To do so, they first computed price parities for the metropolitan areas which the BLS surveys (Aten 2006). They then combined these price parities with county–level data on house prices to produce estimates of price levels by county. Lastly, they aggregated these price parities to the desired level.

We use these data to construct an average cost of living for agricultural workers and non–agricultural workers by state. The execution of this approach is complicated by the fact that the Census obscures the exact metropolitan area of residence for workers in small metropolitan areas or on the edges of some large metropolitan areas to protect their privacy. As a result of this restriction, the residence information from the Population Census falls into one of three categories: (i) the worker lives in an identified metropolitan area; (ii) the worker lives in the rural (i.e., non–metropolitan) area of an identified state; (iii) the worker lives in an unidentified metropolitan area of an identified state. We assign the following price levels to these three residence categories: (i) the price level for the metropolitan area of residence; (ii) the rural price level for the state of residence; (iii) the average price level for the unmatched metropolitan areas for the state of residence. An additional technical complication arises with metropolitan areas that span multiple states. We apportion these metropolitan areas (and their price levels) to the individual states using county–level employment data. We weight using nominal compensation at the county level, as did the original research.

D Labor Shares at the State Level

In this section, we explain in more detail how we estimate the labor shares in agriculture at the state level. The difficulty in doing this is the usual one, that is, some of agriculture value added accrues to proprietors as compensation for the capital and labor inputs that they provide, and so we need to split proprietor's income between the payments to labor and capital. We now describe the detailed steps involved in calculating the labor shares in agriculture at the state level that we use in the text and then compare these labor shares with those that result by applying a different method which is due to Gollin.

To calculate the labor shares at the state level that we use in the text, we impute the different factor payments made in agriculture in each state, and then use the ratio of the imputed labor payments to the imputed total factor payments as our estimate of the state's labor share.

FOR WHICH YEAR?

We begin with the payments made to capital in each state, which we divide into the payments made to land and to physical capital. To impute the payments to land in agriculture for each state, we draw on Turner et al. (2011), who provide for each year the acreage and the dollar rental price of the three types of land used in agriculture: irrigated cropland, non–irrigated cropland, and pastureland. We obtain the rental payments to land by multiplying each state's acreages of the three types of land by their dollar rental prices and summing the results. To impute the payments made to physical capital in agriculture for each state, we use that Turner et al. (2011) also provide the dollar values of the physical capital stock in agriculture for each state, including buildings, machinery and equipment, livestock, and crop inventories. We assume a standard real return of 4% and use a the depreciation rate of 7.51% (which is what the BEA reports for agriculture during the relevant time period) to arrive at a gross rental rate for agricultural physical capital of 11.51 percent. We obtain the rental payments for physical capital by multiplying each state's capital stock by this gross rental rate.

We impute the payments to labor in agriculture in each state as follows. The Census of Agriculture includes the total payments, including benefits, to wage workers at the state level. Therefore, we only need to estimate the payments to self–employed and unpaid workers. We do so using the CPS, which conveniently identifies self–employed and unpaid workers. We could use their reported information on wages, but we have the usual concern that self–reported wages are unreliable for these workers. Therefore, we instead impute their wages. To do this, we first run a wage regression using workers who are employed for wages, controlling for state, age, gender, and education. We then use the results of this regression to impute the hourly wage for self–employed workers. We combine their hourly wage with their hours worked to estimate their total labor payments. We sum these labor payments by state to find the labor payments to self–employed and unpaid workers by state.

We estimate the labor share as the ratio of imputed labor payments to the sum of all imputed factor payments. This ensures that the capital and labor share sum to unity. We note that the sum of imputed payments is not equal to value added, which is not too surprising given the variability in the estimates of value added that we found in the main body of the paper. The resulting estimates of the labor share are low relative to the rest of the economy. For example, the national labor share in agriculture of 0.35. There is also tremendous variation in the labor share across states, ranging from a low of 0.13 in Iowa to a high of 0.81 in New Hampshire.

Given the large variation in labor shares implied by our benchmark method, we provide additional evidence from the method of Gollin (2002). The basic idea of Gollin is to identify the factor payments that unambiguously go to labor and to capital and then split the rest of value added (i.e., proprietors' income) in the same proportions. We use the 2002 Census of Agriculture as our data source, because it provides a more detailed record of factor payments than do other data sources. Some items in the 2002 Census of Agriculture are unambiguously capital or labor payments:

- Labor payments: expenses for "hired farm labor" and "contract labor". The former includes benefits for farm workers.

- Capital payments: "cash rent for land, buildings, and grazing fees", "rent and lease expenses for machinery, equipment, and farm share of vehicle sales", "interest expense", and "landlord's share of total sales" net of "production expenses paid by landlord".

After subtracting these categories, a residual value added remains, which represents payments to proprietors and the profits of corporate farms. Unlike in other sectors, this residual is relatively large. For the median state, 43 percent of value added cannot be directly attributed to capital or labor payments, and in one state the residual is as large as 70 percent. This is what we referred to in the text when we said that Gollin's method has the weakness that in agriculture proprietors' income is a very large share of value added in some states.

Figure 11 shows the relationship between the estimates of labor shares that we used in the text and the estimates produced by Gollin's method. The figure conveys two important points. First, the two series are highly correlated, with a correlation coefficient of 0.89. This suggests substantial agreement between the two methods about which states have more or less labor intensive agriculture. Second, the estimates from Gollin's method display even *more* variation than those from our method, actually ranging from a low of 0.16 in Iowa to a high of 0.86 in Hawaii.

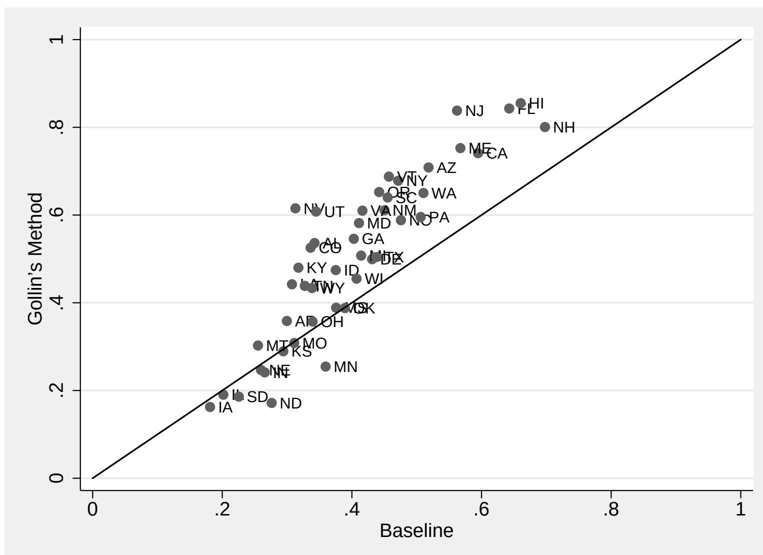

Figure 11: Comparison of Two Estimates of Capital Shares in Agriculture