Why is Measured Productivity so Low in Agriculture?*

Berthold Herrendorf and Todd Schoellman
Arizona State University†

September 30, 2012

Abstract

We document for selected samples of US states and countries that labor productivity and wages are both much lower in agriculture than in the rest of the economy. We then establish that: the productivity gaps are implausibly large compared to the wage gaps; there is a measurement problem in agriculture; the wage gaps are mostly accounted for by the fact that human capital is lower in agriculture. For US states we establish in addition that correcting the measurement problem reduces the productivity gaps by so much that they become consistent with the wage gaps.

Keywords: productivity gaps; wage gaps; measurement problem; human capital.

JEL classification: O1.

*For helpful discussions and suggestions, we thank Roozbeh Hosseini, Moritz Kuhn, David Lagakos, Patrick Minford, Edward Prescott, Richard Rogerson, Ákos Valentinyi, and Gustavo Ventura, as well as seminar participants at Arizona State University, the Canadian Macro Study Group, the Cologne Workshop on Macroeconomics, the Free University of Berlin, the Growth and Development Conference at Toronto, the HKUST Macroeconomics Conference, Mannheim, the Midwest Macroeconomics Meetings, the Meetings of the Society for Economic Dynamics, UC Santa Barbara, and the Universities of Toronto and Western Ontario. 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 **1**

2 Conventional Measurement **4**

2.1 Measurement for 2000 5

2.2 Improved measurement 6

2.2.1 Value added 6

2.2.2 Hours worked 6

2.2.3 Averages by decade 8

3 Comparing Gaps in Productivity and Wages **9**

3.1 A key identity 9

3.2 Measuring wages 11

3.3 Locating the measurement problem 13

4 Re-measuring Agricultural Productivity **14**

4.1 Hours 14

4.2 Farm value added 15

4.3 Proprietors' income 16

5 Human Capital **19**

6 Evidence for Selected Countries **21**

7 Conclusion **27**

A Data Appendix: Sectoral Value Added **32**

A.1 BEA 32

A.2 USDA 32

B Data Appendix: Sectoral Labor Inputs **34**

B.1 Population Census 34

B.2 Current Population Survey 37

C Data Appendix: Proprietors' Income **39**

C.1 Census of Agriculture 39

C.2 March supplement of the CPS 41

D Data Appendix: Consumer Price level 41

1 Introduction

Developing countries have large shares of their labor forces in agriculture although they are much less productive there than in the rest of the economy; see for example, Caselli (2005) and Restuccia, Yang and Zhu (2008). This fact has created a lot of recent interest, because it suggests that it reallocating workers from unproductive agriculture to productive non-agriculture may lead to considerable increases in aggregate GDP per worker. Before one can draw this conclusion, however, one needs to answer the question, What are the reasons for the large sectoral productivity gaps? Recent attempts to provide an answer point to the scale or risk of farming (Adamopoulous and Restuccia 2011, Donovan 2012); to barriers of moving workers or intermediate goods between agriculture and non-agriculture (Restuccia et al. 2008, Herrendorf and Teixeira 2011); to differences in factor endowments (Caselli 2005); to selection of the workers in the two sectors (Lagakos and Waugh 2010); and to home production (Gollin, Parente and Rogerson 2004). Although these contributions have improved our understanding of what is special about agriculture, it is probably fair to say that to date we do not have a conclusive answer to the question posed above. One reason for this is the scarcity of reliable evidence from developing countries, which makes it hard to connect the different theories to the data.

In this paper, we introduce two new ideas to the literature. The first one is that we can learn more about the determinants of productivity gaps from studying agriculture in US states, for which we have a host of well documented data from comparable sources such as the BEA, the CPS, and the USDA. The second new idea is that we can learn more about the determinants of productivity gaps from studying wage gaps, which are closely related to productivity gaps. Calculating wage gaps requires data about earnings and employment by sector. Since we have these data for most US states, we start our analysis by studying US states. Afterwards, we extent our analysis to sample of twelve countries for which we have sufficient data. Our sample includes countries from all stages of development.

As a first pass at the evidence from US states, we calculate the labor productivity by sector for the year 2000, where labor productivity is defined as value added per worker in current dollars. For this first pass, we use data sources that are as similar as possible to those available across countries, namely the BEA's regional accounts (which underly NIPA) for value added and the Population Census for employment. We find that in most US states labor productivity was considerably higher in non-agriculture than in agriculture; the median gap was almost a factor two and the maximum gap was a factor of three. In other words, there are surprisingly large productivity gaps in US states in 2000 that are in the same ballpark as those typically found for poorer countries. One potential issue with this finding is that the employment data from the Population Census records bodies in the first job instead of hours worked, so actual employment by sector will differ from census employment if the hours worked or the relative importance of second jobs differ by sector. The advantage of working with US states compared to poorer countries is that we can address this issue by using the CPS, which reports hours worked by sector in the first two jobs. We use this information to measure sectoral hours during the thirty–year period 1980–2009. To make sure that we have enough observations in agriculture, we report averages by decade (the 1980s, 1990s, 2000s) and we exclude the five states with the smallest agricultural sectors. We find that the productivity gaps get even larger than before, particularly at the right tail where the maximum gap now is 5.7.

Having established that there are large and sustained productivity gaps in US states, we turn to measuring wage gaps between non-agriculture and agriculture. The reason for being interested in wage gaps in the context of this paper is that wages and productivity in a sector are connected through a simple accounting identity: the average wage equals average labor productivity times the labor share (defined as the share of value added paid to labor). This simple accounting identity implies that the gap in wages between non-agriculture and agriculture equals the gap in productivity times the gap in labor shares. We provide substantial evidence that the labor share in non-agriculture is higher than in agriculture, primarily because agriculture is more land intensive than the rest of the economy. In other words, non-agriculture is more productive than agriculture at the same time as which it pays a higher share of value added to labor, implying that the gap in wages between non-agriculture and agriculture must be even larger than the gap in productivity. We calculate wages gaps from the CPS and use the relationship between gaps in wages and productivity as a check on the plausibility of measured productivity gaps. We find that for most states there are sizeable wage gaps, but the wage gaps are considerably *smaller* than the productivity gaps. Since the relationship between wage gaps and productivity gaps is an identity, this means that there must be a measurement problem. To locate the measurement problem, we go a step further and calculate the implied labor shares by sector for each state. We find that while the implied labor shares in non-agriculture are fairly tightly distributed around the standard value of 2/3, the implied labor shares in agriculture are all over the map with many realizations larger than 2/3 and some realizations larger than one. This points to agriculture as the sector with the measurement problem.

Given the detailed and well documented data that are available for US states, we are able to go yet further and show that the measurement problem is in fact with agricultural value added. It turns out that the reason for this is that the BEA seriously under-estimates value added in agriculture along two dimensions: (i) it does not include some factor payments that conceptually belong to agriculture, an example being land rents that it counts as value added in real estate instead of value added in agriculture; (ii) it does not correct sufficiently for under-reporting of proprietors' income in agriculture. We show that making the appropriate corrections for these flaws reduces the measured productivity gaps by so much that they become broadly consistent with wage gaps. The natural last question left to ask then is, What accounts for the measured wage gaps? To answer this question, we use the same wage data from the CPS as above and calculate human capital for non-agriculture and agriculture at the state level in the standard Mincerian way. We find that human capital is much higher in non-agriculture, and that the sectoral differences in human capital account for almost all of the wage gaps. Since the corrected productivity gaps and wage gaps are broadly consistent with each other, this also means that for US states the sectoral differences in human capital account for most of the corrected productivity gaps.

Our results for US states provide an example in which large measured productivity gaps are mostly accounted for by mis-measurement of agricultural productivity and sectoral differences in human capital. This implies that frictions such as barriers to the movement of workers across sectors do not play a quantitatively important role in US states even though measured productivity gaps are large.

Our results are of interest to the development literature to the extent to which they hold also for countries that are at lower stages of development than the US. In the last part of the paper we offer evidence that the key facts from US states about wage gaps continue to hold for a sample of 12 selected countries. We include all countries in our sample for which IPUMS reports data on sectoral earnings and employment from population censuses. This is the case for both rich countries such as Canada, Israel, and the US and poor countries such as Brazil, India, Indonesia, and Mexico. We show that the standard stylized facts about agriculture hold for these countries, that is, the poorer countries tend to have larger employment shares in agriculture and larger productivity gaps between non-agriculture and agriculture; the poorest countries in our sample have around half of their employment in agriculture and very large productivity gaps (the maximum gap is about a factor of five). More importantly, we establish that in this sample of countries there are also large wage gaps, but the wage gaps are considerably smaller than the productivity gaps; the implied labor shares are plausible in non-agriculture and implausible in agriculture, suggesting that there is a measurement problem in agriculture; the wage gaps are mostly accounted for by the fact that workers in non-agriculture have higher human capital than those in agriculture.

Our paper joins a recent literature about sectoral productivity differences across countries, which aims to identify the sectors that make poor countries unproductive; see for example Herrendorf and Valentinyi (2012). A large part of this literature focuses on the productivity gap between non-agriculture and agriculture, because agriculture has most employment in poor countries. Our paper is most closely related to Gollin, Lagakos and Waugh (2011), which also seek to measure and account for gaps between non-agricultural and agricultural productivity. The main difference between the two papers is that Gollin et al. focus on a much larger set of countries including the poorest ones from Africa. While that makes their sample more representative than ours, the earnings data by sector from population censuses that we use here do not exist for most of their countries. With regards to the conclusions, the two papers nicely complement each other along at least two dimensions. First, Gollin et al. (2011) also find that in countries for which reliable estimates exist, sectoral differences in human capital are an important explanatory factor of productivity gaps. Second, Gollin et al. (2011) are unable to account for a sizeable part of the productivity gaps with the standard explanatory factors. Our evidence on wage gaps suggests that these un-accounted parts of productivity gaps may be due to persistent mis-measurement of agricultural productivity.

The remainder of the paper proceeds as follows. Section 2 measures productivity gaps in US states. Section 3 measures wage gaps in US states and compares them to productivity gaps. Section 4 shows that there is severe mis-measurement of agricultural value added in US states. Section 5 establishes that in US states human capital gaps account for most of the productivity gaps. Section 6 extends the previous analysis to other countries. Section 7 concludes. An appendix contains a detailed description of our data sources.

2 Conventional Measurement

In this section, we document that measured productivity in agriculture is lower than in non-agriculture in most US states, and that the difference is often large quantitatively. Productivity is defined as value added in current dollars per worker or per hour. We call the ratio between the productivity in non-agriculture and in agriculture the *productivity gap*. We will first measure productivity gaps using the conventional data sources that are typically available for developing countries. We will then improve the measurement by using better data on hours worked than is typically available for developing countries.

We follow the FAO and define the agricultural sector as the farm industries crop and animal production. This means that for now we do not include forestry, fishing, and hunting in agriculture. Non-agriculture comprises all industries other than agriculture. We also exclude the military from non-agriculture, because for US states we do not have employment data by state for it.

2.1 Measurement for 2000

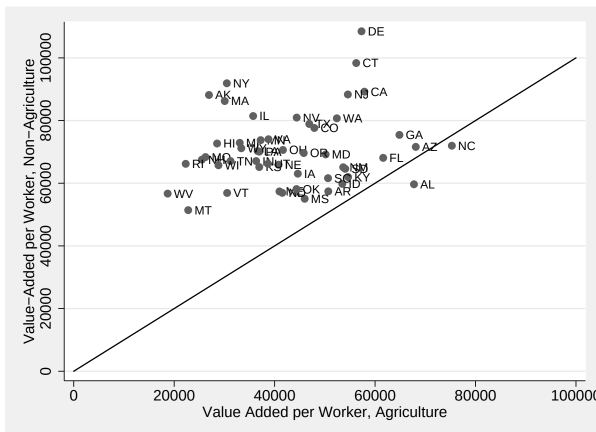

Figure 1: Sectoral Value Added per Worker in 2000

To take a first look at the data, we start with the census year 2000 and data sources for US states that are as similar as possible to those that are available for developing countries. The standard source for agricultural value added is the National Income and Product Accounts (NIPA) as published the United Nations. The standard data source for employment are labor surveys or population censuses. The closest comparable data sources for US states are the BEA's regional accounts (which form the basis of NIPA) and the Population Census.[1]

Figure 1 shows the productivity gaps for 2000. Agricultural value added per worker is depicted on the x-axis and non-agricultural value added per worker is depicted on the y-axis. For almost all states, productivity is higher in non-agriculture than in agriculture. Moreover, the resulting productivity gaps are sizeable: the median gap equals 1.7, the gap at the 90$^{\text{th}}$ percentile equals 2.8, and the maximum gap is larger than 3. It is surprising, at least to us, that there were such large sectoral productivity gaps in some US states in 2000. In the next subsection, we explore whether these gaps are robust to improved measurement.

[1] We obtain the Census numbers from the public–use version made available through Ruggles, Alexander, Genadek, Goeken, Schroeder and Sobek (2010). The Appendix contains a detailed discussion of the data sources and how we construct agriculture and non-agriculture.

2.2 Improved measurement

2.2.1 Value added

The BEA measure of value added by state does not include subsidies and taxes. This is potentially important here because agriculture receives higher than average subsidies.[2] It turns out, however, that it hardly matters for productivity gaps at the median, the 90^{th} percentile, and the maximum whether or not we include subsidies and taxes. The reason for this is that the vast majority of US farm subsidies go to large farms in the few states that have relatively productive agricultural sectors. In contrast, agriculture in relatively unproductive states tends to pay similar taxes as subsidies, implying that the net transfers from the government tend to be small. Therefore, the summary statistics reported above would hardly change if we included subsidies in agricultural value added. More detailed results are available upon request.

2.2.2 Hours worked

Since census employment in the US refers to bodies in the first job, the question arises whether our results change if we use hours worked and take into account multiple jobs. To assess whether this is the case, we measure sectoral hours worked using the monthly data from the *Current Population Survey (CPS)*, which we access through the National Bureau of Economic Research. The CPS is a rotating panel survey administered by the BLS to a sample of households. The survey contains questions about the identity of workers' jobs, about the main job, and about total hours worked. From 1994 onwards, the survey also asks workers in outgoing rotation groups about the identity of their first two jobs and their hours worked at each. Before, the CPS included supplements for the month of May in 1979, 1980, 1989, and 1991 that collected similar information. We use the monthly files in combination with the information from outgoing rotation groups and earlier supplements to estimate the hours worked in agriculture and non-agriculture for each state, month, and year, accounting as well as we can for first and second jobs. The details are available in the Appendix.

To get a sense about the size of the difference between the employment from the Popu-

[2] In its various documentations, the BEA is not very explicit about whether or not it includes net subsidies in sectoral value added. We established that it does not do this at the state level by personal communication with people in the BEA who are familiar with the relevant NIPA procedures. We also compared the BEA numbers for agricultural value added at the state level with value added without subsidies that we constructed from the State–Net–Value–Added Accounts provided by the USDA. We found that the two numbers are very similar.

lation Census and the CPS, we start with the national level. CPS employment is full–time equivalent employment that we calculated on the basis of hours worked in first and second jobs under the assumption that a full–time job amounts to 48 weeks per year and 40 hours per week. We find that CPS employment exceeds Census employment in both sectors. Since the CPS is deemed to be a very reliable source of US employment data, this implies that the Population Census underestimates sectoral employment. We also find that the relative discrepancy between the Population Census and the CPS numbers is considerably larger in agriculture than in non-agriculture. There are two reasons for this. First, the Population Census records only primary jobs, and so it underestimates employment in both sectors. While in principle this could be equally important in both sectors, it turns out that it is more important in agriculture than in non-agriculture because relatively more jobs in agriculture are second jobs. Second, the Population Census is taken during the month of March. As Figure 2 shows, this is more important in agriculture than non-agriculture, because employment in agriculture is very seasonal and March is a month with below average activity in agriculture; in contrast, employment in non-agricultural is hardly seasonal.[3]

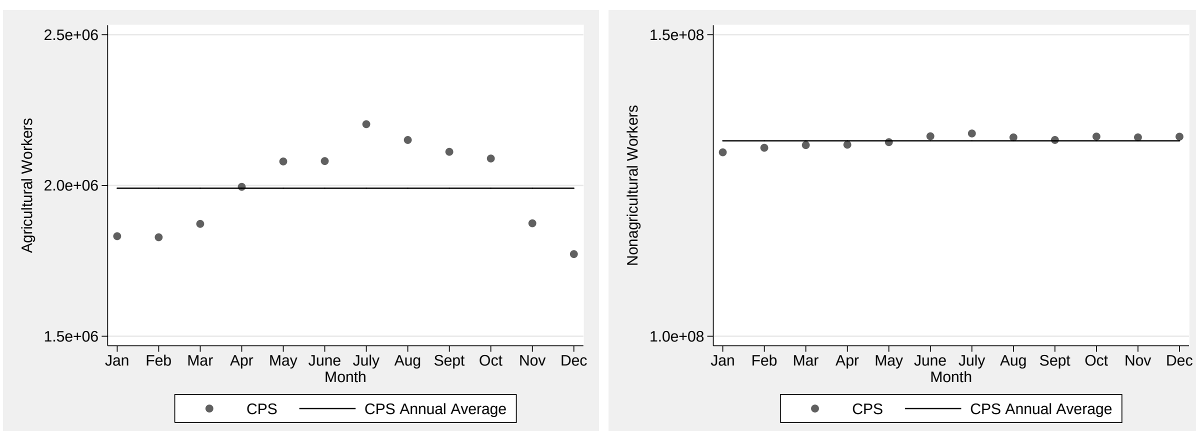

(a) Agriculture

(b) non-agriculture

Figure 2: Sectoral employment during 2000

An additional issue with the employment numbers from the Population Census is that they refer to workers instead of hours worked. If hours per worker are roughly the same in both sectors, then using workers instead of hours worked does not affect the calculations of productivity gaps. The information on hours worked by sector from the CPS allows us to assess whether this is the case for the United States. We find that it is not, and that

[3]Note that hours worked and usual hours worked in agriculture show similarly strong seasonal variation.

workers in agriculture tend to work more hours. Restricting attention to workers with only one job, the average farmer works 40.4 hours whereas the average non-farmer works just 37.2 hours. Not taking the sectoral difference in hours worked into account leads to an underestimation of sectoral productivity gaps.

2.2.3 Averages by decade

Our results for the year 2000 raise the question whether 2000 was an unusually bad year for US agriculture. To address this question, we extend our measurement to 1980–2009. We group the data in three non-overlapping ten–year bins: 1980–1989, 1990–1999, and 2000–2009, which we refer to as the 1980s, 1990s, and 2000s. All numbers that we will report are the ten–year averages of the sectoral productivity gaps by state within the respective bin.[4] While using ten–year bins increases the number of observations underlying each statistic that we calculate, the CPS still contains only a few agricultural workers in states with 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 decade. This criterion leads us to exclude Alaska, Connecticut, Massachusetts, Rhode Island, and West Virginia.

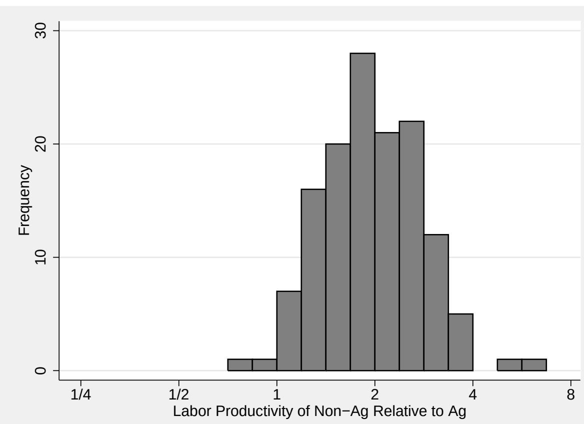

(a) Histogram

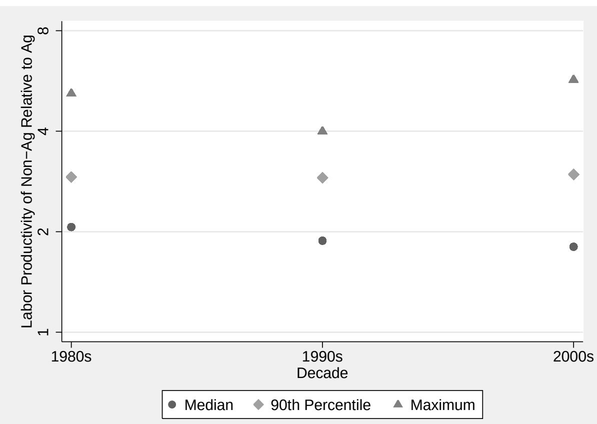

(b) Panel

Figure 3: Improved Measurement of Labor Productivity Gaps 1980–2009

Panel (a) of Figure 3 shows the new measures of productivity gaps in the form of a histogram, showing that with the improved measurement and thirty years of data, the stylized

[4]Note that although the CPS started before 1980, only since 1978 does it have information for each individual state, which is crucial here. We start in 1980 because this allows us to form natural ten–year bins.

Table 1: Productivity Gaps with CPS Hours 1980–2009

Median	1.9
90^{th} Percentile	3.0
Maximum	5.7

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

fact of Subsection 2.1 survives; in most states and years, there is a sizeable productivity gap between non-agriculture and agriculture. Panel (b) of Figure 3 shows that the productivity gaps do not decline over time.[5] This, and the fact that we average over ten–year bins, addresses the concern that our initial results reflected a bad harvest during the year 2000. Table 1 gives the summary statistics for the productivity gaps: the median gap is 1.9, the gap at the 90^{th} percentile is 3.0, and the maximum gap is 5.7. These gaps are larger than the gaps for 2000, although we have excluded five states from the sample. It is remarkable that the summary statistics for US states come out very 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^{th} percentile is 8.8.[6]

3 Comparing Gaps in Productivity and Wages

In this section, we bring to bear additional evidence from sectoral wages. We first derive a key identity that connects productivity gaps with wage gaps. We then measure wage gaps and show that, in comparison, the measured productivity gaps are too large to be plausible.

3.1 A key identity

We use the following notation: Y denotes value added, L hours worked, W the wage per hour (i.e. earnings divided by hours worked), and LS the labor share (i.e., the payments to labor divided by value added). Both Y and W are measured in terms of current dollars. The indexes a and n indicate agriculture and non-agriculture.

[5] It is also the case that the identities of the states with the largest productivity gaps remains the same across decades.

[6] Some readers might wonder whether across US states productivity gaps show the same relationship with state GDP per capita as they do across countries, that is, whether poorer states have higher productivity gaps than richer states. We find that across US states there is no strong relationship between productivity gaps and state GDP per capita.

To derive a relationship between productivity gaps and wage gaps, we recall the definition of the labor share:

$$LS \equiv \frac{WL}{Y} = \frac{W}{Y/L}. \tag{1}$$

We stress that this equation is an accounting identity that must hold without further assumptions on the market structure or the technology. Dividing the identities for non-agriculture and agriculture by each other and rearranging, we obtain a new identity that connects gaps in productivity and wages:

$$\text{Gap}(Y/L) = \text{Gap}(W) \cdot \text{Gap}(LS)^{-1}, \tag{2}$$

where the gap of variable X is defined as:

$$\text{Gap}(X) \equiv \frac{X_n}{X_a}.$$

Identity (2) shows that the productivity gap equals the product of the wage gap and the inverse of the labor share gap. One implication is that the relative sizes of the gaps in productivity and wages depend on the relative sizes of the labor shares in non-agriculture and agriculture. Another implication is that one should not view $\text{Gap}(Y/L) = 1$ as the natural benchmark, because the right–hand side does not equal one in general.

A natural conjecture is that agriculture is less labor intensive than non-agriculture, because agriculture is more land intensive and land is one form of capital. A sizeable body of evidence confirms this conjecture for the US. For example, using the value added data from the BEA and the methodology of Gollin (2002), we find a value of 0.44 for the average labor share in US agriculture during 1980–2009. There is some concern that the methodology of Gollin may not work well for agriculture, because it effectively calculates the labor share for the part of value added that is not produced by proprietors.[7] This part of value added is relatively small when there are many proprietorships like in agriculture. Two additional pieces of evidence show that Gollin's method leads to reasonable estimates even in agriculture. First, Mundlak (2005) pointed out that share-cropping arrangements allocate around half of the harvests to the share cropper. Since sharecropper tend to use some capital they own, this implies a labor share in agriculture that is smaller than 0.5. Second, (Griliches 1964) directly estimated a Cobb–Douglas production function for

[7] The reason for ignoring the value added produced by proprietors is that it is unclear how to split it between labor and capital inputs.

agriculture for the 1950s and found values for the labor share in the range of 0.4–0.5. (Kislev and Peterson 1987) confirm this result for the 1980s.

Since agricultural value added amounts to only a few percentage points of US GDP, non-agriculture comprises almost the whole economy. Hence, we can use the standard aggregate labor share of 2/3 for non-agriculture. With labor shares of 0.44 in agriculture and 2/3 in non-agriculture, (2) implies that

$$\text{Gap}(Y/L) = \text{Gap}(W) \cdot 0.7 < \text{Gap}(W). \tag{3}$$

We view (3) as staking out a reasonable ballpark for the measured wage gaps in comparison to the measured productivity gaps. Next, we will measure wage gaps by state and check whether they are in the ballpark.

3.2 Measuring wages

Figure 4: Wage Profiles by Sector, 2000 CPS

(a) Years of Schooling (b) Potential Experience

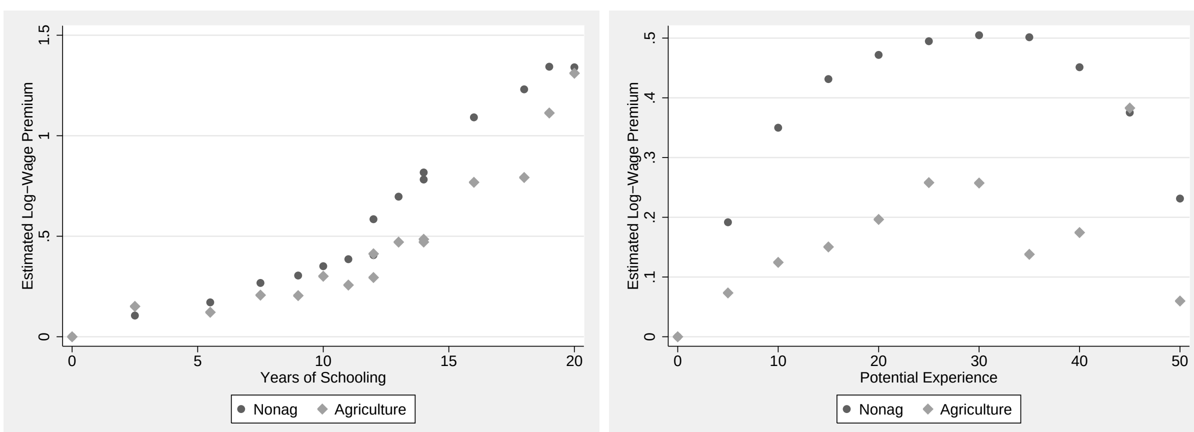

The CPS Matched Outgoing Rotation Groups contain information on hourly nominal wages, age, education, gender, state and sector. We use this information to calculate average wages in current dollars in each sector–state pair. While doing this is straightforward in principle, an issue arises because self-employed, proprietors, and non-wage workers do not report hourly wages. This raises a problem because the employment share of these individuals is fairly large in agriculture; wage information is missing for 61, 58, and 53 percent of all individuals working in agriculture during our three decades, respectively. In

Figure 5: Comparing Wage Gaps to Value Added Gaps in US States

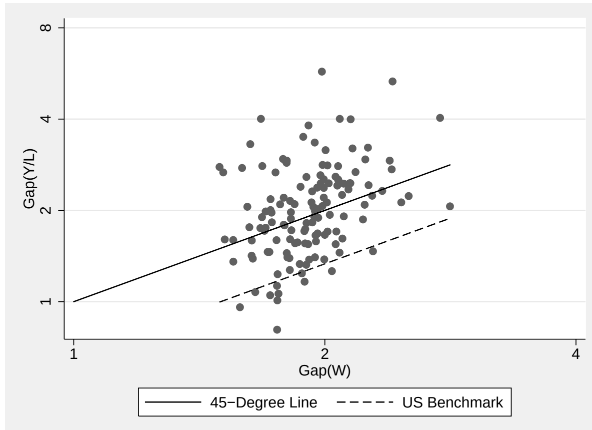

contrast, in non-agriculture wage information is missing only for 11 percent of individuals in each decade. Moreover, it turns out that the individuals with missing wage information differ systematically from those who report wages; for example, proprietors have more years of schooling and more experience than wage workers, and the differences between the two groups are particularly pronounced in agriculture.

To take these differences in observable characteristics between individuals with and without reported wages into account, we impute the missing wage information under the assumptions that self-employed, proprietors, and non-wage workers make the same wage as wage workers with the their characteristics. Specifically, we run regressions of log hourly–wages on state fixed effects and age, education, and gender for the individuals who report this information; we then multiply the observable characteristics of the individuals with missing wage information with the estimated regression coefficients. We choose to run the log–wage regressions separately for each sector and decade, which allows us to capture differences in the wage structure and in the return to observed factors over time and across sectors. One reason why this may matter is that the market return to skills may have risen over time or it may be higher in one sector than in the other. Figure 4 suggests that the latter is indeed the case over the period 1980–2009; the rates of return to schooling in agriculture are lower than in non-agriculture and the profile of potential experience is flatter in agriculture.

Figure 5 plots the productivity gaps against the wage gaps and Table 2 reports the summary statistics. We can see that there are relatively large differences between wages in non-agriculture and agriculture at the state level. Moreover the summary statistics of the wage and the productivity distributions are surprisingly similar. In fact, for the states

Table 2: Wage Gaps

	Gap(Y/L)	Gap(W)	Gap$(LS)^{-1}$
Median	1.9	1.9	1.0
90^{th} Percentile	3.0	2.2	1.7
Maximum	5.7	2.8	2.9

at the median of each distribution, the gaps in nominal wages and human capital are both equal to 1.9. This means that in the median state, an average worker in non-agriculture has almost twice the nominal wage of an average worker agriculture. The gaps at the 90^{th} percentile and the maximum are even larger.[8]

Figure 5 and Table 2 also show that the wages are not in the ballpark staked out by (3): for many states they come out smaller or equal than the productivity gaps, whereas (3) suggests that the opposite should hold. Moreover, in almost all states the implied inverse of the labor share gap is considerably larger than the ballpark figure 0.7. Since the relationship (2) between wage and productivity gaps is an identity that must hold, this implies that there must be a measurement problem either with wages, productivity, or both.

3.3 Locating the measurement problem

While the previous evidence implies that there must be a measurement problem, it does not imply in which sector the problem is. We now use additional information to show that the measurement problem is in agriculture. We do this by calculating the labor shares at the sectoral level that follow when we plug the measured wages and productivities into identity (1). The left panel of 6 shows the implied labor shares for non-agriculture. The decadal averages fall into the relatively tight range $0.59 - 0.64$, which is roughly consistent with the standard values for labor shares of $0.6 - 0.7$. We conclude that the implied labor shares for non-agriculture are plausible.

The right panel of 6 shows the implied labor shares in agriculture. They are very dispersed. Moreover, the average by decade is $0.60 - 0.61$, which is considerably larger than the standard values for the labor shares for agriculture; for example, as we saw above, according to BEA data the average labor share during 1980–2009 was 0.44. Moreover, quite a few observations on the right panel are larger than 1. Given that the underlying

[8]Note that in Table 2 the states for which a given summary statistic of nominal wages are reported may not be the same.

Figure 6: Implied Labor Shares

(a) Non-agriculture (b) Agriculture

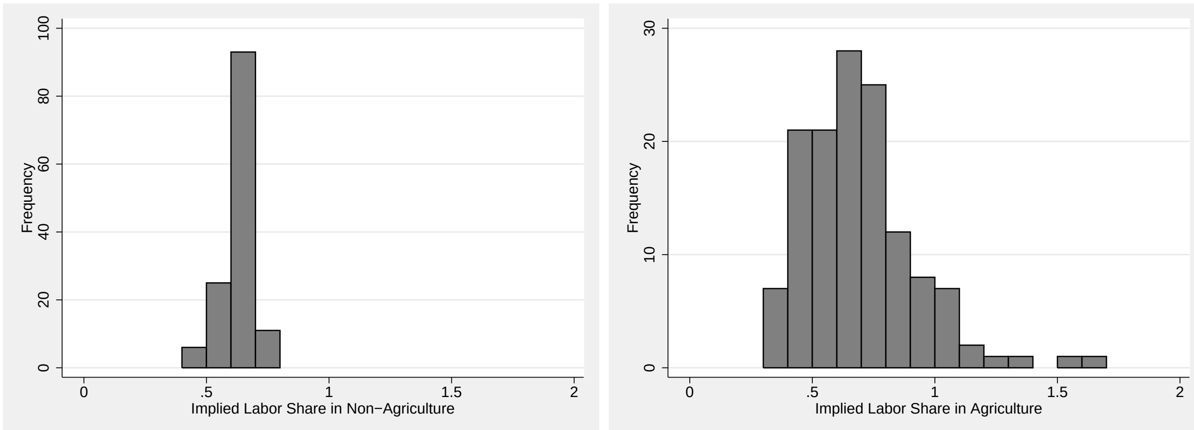

wages and productivities are averages over decades and states, this is not plausible. We conclude that the implied labor shares in agriculture are implausibly large and that the measurement problem is in agriculture. Reinspecting (1), we also conclude that given that the measurement problem is in agriculture, either agricultural wages must be overestimated and/or agricultural productivity must be underestimated. If agricultural productivity is underestimated, then this may be due to agricultural value added being underestimated, agricultural hours being overestimated, or both.

4 Re-measuring Agricultural Productivity

In this section, we establish that agricultural productivity is mis-measured in US states. Showing this requires to show that either the numerator or the denominator of agricultural productivity is mis-measured.

4.1 Hours

To assess whether the denominator is mis-measured, we compare hours worked in the CPS with those in American Time Use Survey (ATUS). ATUS data are collected from time use diaries in which the participants record their activities by minute for a 24–hour period. ATUS data are deemed very reliable. We focus on the period 2003–2010 during which the ATUS and the CPS overlap. Given that ATUS does not have many observations in agriculture, we only calculate hours at the national level. We find that there are only small differences between ATUS and CPS estimates of hours worked: while ATUS hours are 9.5% lower than CPS hours in agriculture, they are 8.4% lower in non-agriculture.[9] This suggests that in the US mis-measurement of hours worked is not an important factor behind the implausibly large productivity gaps that we found above. In the following subsections, we will establish that instead mis-measurement of agricultural value added is the key factor behind them.

4.2 Farm value added

Above, we have used farm value added from the BEA's regional accounts. Although these data underlie the construction of NIPA, it turns out that they have a crucial shortcoming when it comes to measuring agricultural value added at the state level. The shortcoming comes from the fact that the BEA follows the System of National Accounts (SNA) and views agricultural value added as the value added that is produced by *farmers*, i.e., persons who operate a farm or who are employed on a farm, instead of the value added that is produced on *farms*. 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 generated on farms, and conceptually they belong to the value added produced on farms. However, the BEA does not report them as part of agricultural value added because they do not lead to income of farmers. Instead it reports the payments to farm contractors as value added in agricultural services and the payments to land owners as value added in real estate.

To see how much this shortcoming matters for measuring productivity gaps, we construct a new measure of value added that includes all factor payments generated on farms irrespective of who they accrue to. To this end, we use 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. We obtain non-agricultural value added as the difference between state GDP and state agricultural value added.

Figure 7 plots on the x-axis the agricultural value added from the BEA, which we

[9] Note that the fact that ATUS hours are smaller than CPS hours is expected because ATUS uses a strict definition of time spent at work. For example, ATUS does not count time spent at business meals or commuting, while respondents in the CPS may implicitly include such time in their responses.

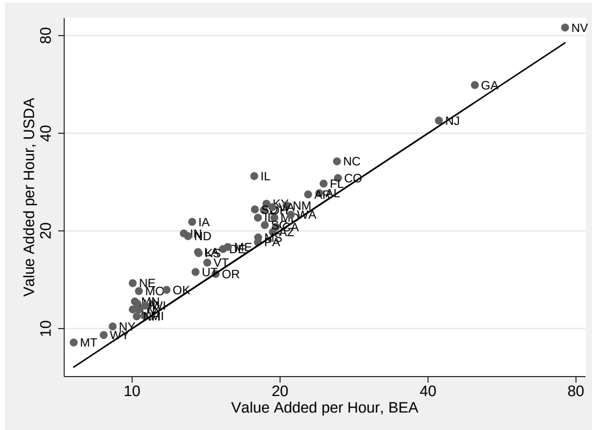

Figure 7: Agricultural Value Added in 2000 from Different Data Sources

used above, and on the y-axis the agricultural value from the USDA, which includes all factor payments. All value added numbers in the figure are for the year 2000 and per CPS hour worked. As expected, most observations are above the 45 degree line, that is, value added per hour worked based on the USDA is larger than value added per hour worked based on the BEA. Moreover, for states with low agricultural productivity the difference can be large quantitatively. This suggests that using the BEA numbers for value added leads to a downward bias in the calculation of productivity in agriculture, and thereby biases upwards the measured productivity gaps between non-agriculture and agriculture. Since BEA value added is similar to the standard data sources that are available for developing countries, this suggests that part of the large productivity gaps that people find for developing countries may be the result of mis-measurement due to the practise of counting agricultural value added that is not generated by farmers in non-agriculture. In the next subsection, we identify a further source of mis-measurement of agricultural value added, which arises because proprietors tend to under-report their income severely.

4.3 Proprietors' income

The Internal Revenue Service (IRS) periodically conducts tax audits to assess the degree of tax compliance. These audits find that proprietors severely under-report their income. Table 3 lists the ratio of actual proprietors' income (as determined by tax audits) to reported proprietors' income for the two studies that fall into our period or investigation, Internal Revenue Service (1996) and Internal Revenue Service (2007). The estimates range from 1.4 to 3.6, suggesting under-reporting on a massive scale. Moreover, the degree of under-

Table 3: Actual divided by Reported Proprietors' Income

	Non-farm	Farm
1980s	$1.4 - 1.5$	$1.4 - 1.5$
2001	2.3	3.6

reporting appears to be more severe in farming than in non-farming. Potentially this is a serious issue for calculations of agricultural productivity because proprietorships are very common in farming.

In non-agriculture, the BEA reacts to the evidence of under-reporting by taking the IRS estimates into account. As a result, it roughly doubles the reported income of proprietors. In agriculture, the BEA follows the production approach to estimate value added as the difference between revenues and expenses. Since the production method does not directly feature the income of proprietors, it is not clear in principal whether the BEA should make any adjustment for the under-reporting of farm proprietors' income.[10] It turns out, however, that in practise the BEA should adjust for the under-reporting of farm proprietors' income. The reason for this is that the value added of farm proprietors under the production approach only modestly exceeds the reported income of farm proprietors. We establish this for 2002 and 2007 by comparing proprietors' income calculated from the March supplement to the Current Population Survey with proprietors' income calculated from the Census of Agriculture. We find that at the national level the ratios of the value added of farm proprietors under the production approach and the reported income of farm proprietors equal 1.1 and 1.2 in 2002 and 2007, whereas the IRS estimates indicate that a much larger correction is warranted.

If under-reporting of proprietors' income in agriculture is part of the reason why we have such large productivity gaps in some states, then states with a larger proprietors' income share in agricultural value added should have larger productivity gaps. Figure 8 shows that this is indeed the case (the regression coefficient is significantly different from zero at the 99% significance level). This suggests that under-reporting of proprietors' income should account for part of the large measured productivity gaps. It is hard to be more specific though. To begin with, over the years the IRS studies found fairly different degrees of under-reporting. Moreover, we have no information on whether the degree of under-reporting in agriculture is uniform or varies by state. This is a possibility because the type of agriculture varies across US states. For example, coastal states like California and Florida specialize in

[10] The preceding paragraph draws on Bureau of Economic Analysis (2009), which contains for details on how the BEA actually constructs state–industry value added.

Figure 8: Proprietors' Share in Agricultural Value Added versus Productivity Gaps (averages over 1980–2009)

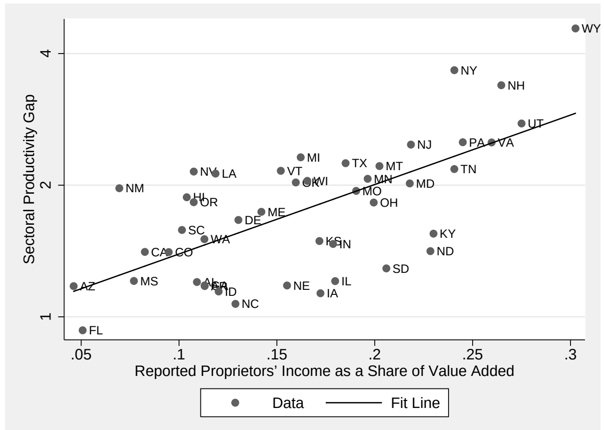

fruit and vegetable production, the Great Plains specialize in grain production, and Texas and the Western states specialize in animal production. Since these types of agriculture have different production structures, e.g. in terms of the degree of mechanization and the reliance on intermediate inputs, the scope for under-reporting of proprietors' income could well differ.

We start by experimenting with the a range $40 - -260\%$ under the assumption that the adjustment factor is the same in all states. This means that our results on the importance of under-reporting of proprietors' income should be interpreted with some care. The productivity gaps after the corrections for missing value added and for under-reporting of proprietors are listed in Table 4 for three different correction factors: 1.5 is the average over the range $1.4 - 1.5$, whereas 2.3 and 3.6 are the suggested adjustment factors for non-farming and farming in 2001. We have highlighted the third column, because it adjusts farm proprietors' income by the same adjustment factor of 2.3 that the BEA uses for non-farm proprietors' income. Intuitively, this column asks how far we can get simply by performing the BEA's adjustment from the non-farm sector also in the farm sector. It is reassuring that for the median state the third column produces a corrected productivity gap of 1.3, which is exactly the value required to resolve the productivity puzzle. To see why, recall that the wage gap for the median state was 1.9, so with a corrected productivity gap of 1.3 the ratio of the two gaps equals $1.3/1.9 = 0.7$, which is what equation (3) suggested it should be. While this works at the median, at the maximum the corrected productivity gaps are still larger than the wage gaps. However, this is not necessarily implausible anymore, because the differences between the two gaps are largely reduced and it is perfectly possible that in

Table 4: Corrected Productivity Gaps 1980–2009

Correct. Factor	1.5	**2.3**	3.6
Median	1.5	**1.3**	1.1
90^{th} Percentile	2.5	**2.0**	1.7
Maximum	4.8	**3.6**	2.9

.

some US states agriculture is more labor intensive than non-agriculture, particularly when the agriculture focuses on fruit and vegetable production.

In sum, we have established two reasons for why the BEA under-estimates agricultural value added at the level of US states: the SNA classifies some of the value added generated on farms in other industries; the BEA does not properly adjust for under-reporting of proprietors' income in agriculture. This implies that the productivity gaps between non-agriculture and agriculture that are based on BEA value added data are artificially large. Caselli and Coleman (1998), which is the unpublished appendix to Caselli and Coleman (2001), reports a related fact: At the US national level during 1940–1990, nominal wages from the Population Census imply a larger reduction in the wage gap between non-agriculture and agriculture than nominal wages from NIPA. Moreover, it turns out that nominal wages in non-agriculture from both data sources are similar and that the difference is due to the behavior of the wages in agriculture. The authors argue that the explanation is that NIPA under-estimates agricultural wages. This is consistent with our finding that NIPA under-measures agricultural value added during the later period 1980–2009 that we study here.[11]

5 Human Capital

Given that the corrected productivity gaps are broadly in line with wage gaps, the question remains what accounts for wage gaps. In this section, we will show that human capital gaps account for most of the wage gaps. To construct human capital, we follow the standard Mincerian practice and use the estimated slope coefficients from our log–wage regressions. In particular, we use the coefficients on schooling, experience, and gender for each sector to

[11]Note that the wage gaps that Caselli and Coleman report are not strictly comparable with our wage gaps. The reason for this is that they compute earnings per worker whereas we compute wages per hour. Since agricultural workers tend to work more hours than the rest of workers, our labor productivity in agriculture is smaller and our wage gaps are larger than theirs.

Table 5: Gaps in Nominal Wages and Human Capital

	Gap(W)	Gap(H)
Median	1.9	1.9
90$^{\text{th}}$ Percentile	2.2	2.1
Maximum	2.8	2.1

construct human capital for this sector given the observed values of schooling, experience, and gender of all individuals who work in it. The reasoning behind this definition is that human capital corresponds to the differences in workers' innate characteristics, evaluated at the observed market rate of return. In contrast, the intercepts and fixed effects measure the determinants of wages that are independent of workers' characteristics (an example would be that unions in some parts of non-agriculture increased the wage above the competitive wage).

Table 5 reports the summary statistics for the gaps human capital; in the second column we repeat the gaps in wages from above for comparison. The gaps in wages and in human capital are surprisingly close to each other. This suggests that in US states most of the wage gaps between non-agriculture and agriculture are accounted for by the fact that the workers in non-agriculture have higher human capital.

The part of the wage gaps that is not accounted for human capital comes from other sectoral differences. The most obvious one is cost–of–living differences between rural areas and cities. To explore the importance of cost–of–living differences, we measure the consumer price level in each location as the average price level that the workers of that location pay for their consumption basket. To calculate this price level for each worker, we need to know the sector in which he works and the price level where he lives. The information about the sector and the residence comes from the Population Censuses.[12] The information about the price level in the residence comes from recent research by Aten (2006), Aten and D'Souza (2008), and Aten (2008), who constructed consumer price levels for 363 metropolitan areas as well as for the rural area of each state for the year 2006.[13] We combine this information with the residence information to obtain a price level for each worker. We then average

[12]After 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.

[13]Note that the CPI is not useful in the current context, because it covers only the major metropolitan areas and is normalized to 100 for every city in the base year, making level comparisons impossible.

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

We find that across the entire United States the gap in the consumer price level between metropolitan and non-metropolitan areas is 1.4. While this gap is fairly large, a considerable part of it comes from price variation across states. Moreover, for the current purpose, we have to take into account that 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 consumer price levels between non-agricultural and agricultural workers are sizeable (around 1.3) only in the few states that have large, expensive urban areas, notably Illinois, New York, Texas, and Virginia. In contrast, in the median state the gap in the consumer price level between non-agriculture and agriculture merely equals 1.1, which is negligible quantitatively compared to the gaps in nominal wages and human capital that we found above.

6 Evidence for Selected Countries

So far, we have focused our analysis on the detailed data from US states that are typically not available for developing countries. This raises the obvious question to what extent our results apply outside of the US. There are good reasons to believe that they should. To begin with, the procedures of the SNA are followed by most countries and under-reporting of proprietors' income is a widespread problem. Moreover, the fact that even the BEA does not correctly estimate agricultural productivity suggests that this is inherently difficult to do compared to the other sectors of the economy. One problem comes from the fact that agriculture has more proprietors than the rest of the economy. In poor countries, this problem is aggravated by the fact that many proprietors are subsistence farmers who do not trade most of their production in markets. Another problem with measuring agricultural value added is that agricultural production takes place in thinly populated areas that are often far from urban areas, which makes hiding output easier in agriculture than in the rest of the economy.

In this section, we leave US states aside and show that, indeed, most of our results hold also for a sample of countries for which we have information about population censuses with the detailed data required for our analysis. In particular, we will show that in these countries: (i) there are large productivity gaps and wage gaps between agriculture and non-agriculture; (ii) the productivity gaps are implausibly large compared to the wage

Table 6: Summary statistics for characteristics of sample countries

	GDP pc Gap (US rel. to country)	Agric. Empl. Share (in %)	Prod. Gap (non-ag. rel. to ag)
Median	4.5	17	2.6
90$^{\text{th}}$ Percentile	11	44	4.3
Maximum	22	62	4.4

gaps; (iii) there is a measurement problem in agriculture; (iv) the wage gaps are mostly accounted for by the fact that human capital is lower in agriculture. To enter into our sample, a country needs to satisfy two criteria: for at least one year during the period 1990–2009 the UN data base reports NIPA information on value added in agriculture as a share of GDP and IPUMS reports employment by sector and earned income from population censuses. We insist on information from population censuses because we want data sets that are large and representative. The following country–year pairs satisfy these criteria: Brazil (1991,2000); Canada (1991,2001); India (1993,1999); Indonesia (1995); Israel (1995); Jamaica (1991,2001); Mexico (1990,2000); Panama (1990,2000); Puerto Rico (1990,2000); Uruguay (2006); United States (1990,2000); Venezuela (1990,2001).

Table 6 reports several summary statistics of our sample of countries: at the median, the gap in GDP per capita (in international prices) with the US is a factor of 4.5, agriculture has an employment share of 17 percent, and the productivity gap between non-agriculture and agriculture is 2.6; at the maximum, these statistics grow considerably to a GDP gap of 22, an employment share of agriculture of 62 percent, and a productivity gap of 4.4. The correlation in our sample between GDP per capita relative to the US and the employment share in agriculture is -0.71 and the correlation between GDP per capita relative to the US and the productivity gaps is -0.65. In addition to rich countries like Canada and the US, our sample also comprises relatively poor countries that have large workforces in the sector in which they are very unproductive, agriculture. India and Indonesia are two examples. One shortcoming of our sample is that it does not contain African countries which have even larger GDP gaps relative to the US than a factor 22. The reason for this is that IPUMS does not report sufficiently detailed census information about wages by sector for African countries.[14]

Having established that there are large productivity gaps between non-agriculture and

[14] South Africa has the best information among all African countries, but that information is still not sufficient to conduct our exercise.

agriculture in our sample of countries, we now conduct the same analysis as for US states to the extent possible. We start by calculating wage gaps. We follow the same principles as before to the extent possible. In particular, to calculate wages we use only information from individuals who are at least 10 years old and have valid information for our key variables: sex, education, earnings, employment by sector and so on. We are again missing wages for proprietors and some workers. As before, we use wage regressions to impute the missing wages. The characteristics that we include in these regressions are the geographic region, education, potential experience, and sex. As before, we run the wage regressions separately for agriculture and non-agriculture.

We face two challenges when imputing wage information for our sample countries. First, in some countries a sizeable fraction of workers list their status as unpaid workers. We take these reports seriously and enter a zero wage income for all such workers. Since agriculture has a higher fraction of unpaid workers than non-agriculture, this may bias our calculated wage gaps upwards. Since we will find that the calculated wage gaps are still too small compared to the productivity gaps, this way of proceeding is conservative and does not invalidate our conclusions. Second, although we have information on earnings from population censuses for all countries in our sample, the level of detail of the employment information differs across countries. In some countries (including the US), we have information on hours worked so that we can calculate the hourly wage for wage workers. We then use that wage in our imputations of the annual earnings of proprietors (which equals the hourly wage of a wage worker with the same characteristics as the proprietor times the number of hours worked by the proprietor). In some other countries, we have less detailed information but can calculate daily or weekly wages. Using these wages for the imputation amounts to assuming that proprietors work the same number of daily or weekly hours as wage workers. In yet some other countries, we know nothing about hours worked so that we are forced to work with monthly income. Using these wages for the imputation amounts to assuming that proprietors work the same number of monthly hours as wage workers. We emphasize that all three methods lead to an estimate of the annual labor income of proprietors. The reason why the first method is preferred to the second one and the second method is preferred to the third one is that the first method needs fewer assumptions than the second one and the second method needs fewer assumptions than the third one. We also emphasize that while not taking into account differences in hours worked between proprietors and wage workers may bias our calculations of the wage gaps, the difference in hours worked between wage workers and proprietors would have to be very large to overturn our results. For the US, we don't find anything close to such large differences in hours worked

Figure 9: Comparing Wage Gaps to Value Added Gaps in US States

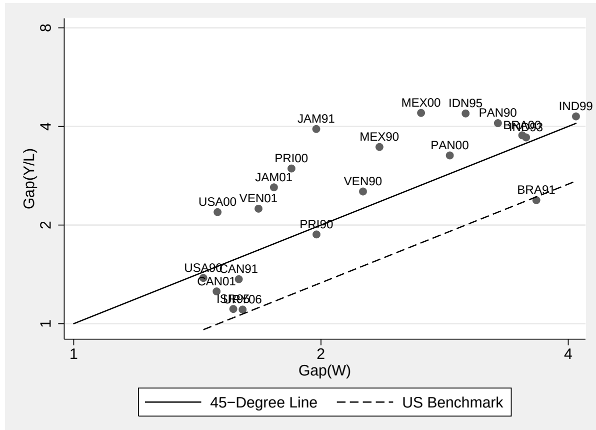

Table 7: Summary statistics for gaps in sample countries

	Gap(Y/L)	Gap(W)	Implied Gap(LS)$^{-1}$
Median	2.6	2.0	1.1
90$^{\text{th}}$ Percentile	4.3	3.6	1.6
Maximum	4.4	4.1	2.0

between proprietors and wage workers.

Next, we compare the gaps in wages and productivity with each other. As above, the relative sizes of the labor shares in agriculture and non-agriculture are crucial in this context. Above, we worked with labor shares in agriculture and non-agriculture equal to 0.44 and 0.66, respectively, and an implied ratio of the labor shares in agriculture to non-agriculture of 0.7. There is a substantial evidence that similar numbers apply also across countries. Thus, we should see larger wage gaps than productivity gaps also across countries.

A first piece of evidence comes from the classic study by Hayami and Ruttan (1970), which found for a sample of thirty eight countries that depending on the estimation method the average agricultural labor share falls into the range 0.34–0.49. Subsequent studies have confirmed that, if anything, the average agricultural labor share is smaller; see for example Fulginiti and Perrin (1993) and Craig, Pardey and Roseboom (1997). A second piece of evidence is that also outside of the US, share–cropping arrangements allocate around half of the harvests to the land owner, which leaves at most the other half for labor (Mundlak 2005). A third piece of evidence comes from specific studies about the countries in our sample: Echevarria (1998) found an agricultural labor share of 0.41 for Canada during 1971–93;

Schultz (1964) found 0.4 for India during 1918–19; Mundlak, Larson and Butzer (2002) found less than 1/3 for Indonesia during 1980–98; Fishelson (1974) found 0.44 for Israel in the late 1960s. We conclude that the existing evidence for our sample of countries implies that agricultural labor shares are smaller than 1/2. Since the existing evidence for aggregate labor shares that they equal 2/3 on average and are not correlated with GDP per capita, this must mean that on average the labor shares in non-agricultural are larger than the labor shares in agriculture.

One might think that the evidence of a smaller labor share in agriculture contradicts Gollin's evidence that the average labor share is uncorrelated with GDP per capita. After all, poorer countries tend to have larger agricultural sectors, and so they should have smaller aggregate labor shares, or should they? There are two reasons why this way of thinking is flawed. First, Gollin's method calculates the aggregate labor share as the share paid to labor in the part of value added that is not produced by proprietors. Since in poor countries agriculture is dominated by proprietors (subsistence farmers to be precise), this means that for poor countries Gollin's method effectively calculates the labor share for non-agriculture, instead of for the aggregate economy. Second, even if all countries had the same labor share in agriculture as the US, they could still have the same aggregate labor share as the US. For this to work out, richer countries would have to have smaller labor shares in non-agriculture than poorer countries. Since the share of agricultural value added in total GDP is relatively small even in poor countries, the required variation in the labor share of non-agriculture would remain modest.[15]

Figure 9 plots and Table 7 reports the results for wage gaps in comparison to productivity gaps. We can see that again the productivity gaps exceed the wage gaps for most country–year pairs. Moreover, at the maximum the productivity gap comes out as twice the wage gap. To reconcile this that with identity (2), we would have to believe that agriculture is more labor intensive than non-agriculture in most countries and that in one country it is twice as labor intensive. This is implausible and suggests that, as before, there is a measurement problem.

The next step is to show that, as before, the measurement problem is located in agriculture. Since the population censuses do not contain information about non-wage income (like benefits), we cannot follow the same steps as for US states and directly calculate the implied sectoral labor shares by combining census information with information about benefits. We therefore go down a different path and calculate the labor shares implied by

[15] In our sample of countries, the share of agricultural value added in total value added remains below 1/3 in all countries.

Table 8: Summary statistics for implied labor shares in sample countries

	LS_n	LS_a
Median	0.67	0.75
90$^{\text{th}}$ Percentile	0.67	1.08
Maximum	0.70	1.25

productivity gaps, wage gaps and the identity:

$$LS = (Y_a/Y)LS_a + (Y_n/Y)LS_n.$$

Solving for the labor shares in agriculture and non-agriculture, we find:

$$LS_a = \frac{LS}{(Y_a/Y) + (Y_n/Y)(LS_n/LS_a)},$$
$$LS_n = \frac{LS}{(Y_a/Y)(LS_a/LS_n) + (Y_n/Y)}.$$

For the right-hand variables of a particular country, we use the standard value of 2/3 for the aggregate labor share LS, the NIPA values for the sectoral value added shares Y_i/Y, and the ratios implied by the previous analysis for the labor–share gaps LS_n/LS_a. Table 8 reports the summary statistics of the results. We can see that the implied labor shares in non-agriculture are close to the standard aggregate value of 2/3, which is plausible. In contrast, the implied labor share in agriculture at the median is already larger than in non-agriculture and at the 90^{th} percentile it considerably exceeds 1. This is not plausible. In other words, the conclusion is similar to that for US states: there is a measurement problem in agriculture also for the sample of countries. Unfortunately, for our sample of countries we do not have sufficient information to establish which exact nature the measurement problem takes, so we will have to skip this step of the previous analysis.

The last step is to establish, that, again as before, the wage gaps are mostly accounted for by the fact that human capital is higher in non-agriculture than in agriculture. Following the same steps as for US states, we calculate human capital by sector. Table 9 shows that the results are very similar to those for US states: there are large gaps in human capital; the sizes of the gaps in wages and human capital are surprisingly close to each other.

Table 9: **Summary statistics human capital in sample countries**

	Gap(W)	Gap(H)
Median	2.0	2.0
90$^{\text{th}}$ Percentile	3.6	3.3
Maximum	4.1	4.3

7 Conclusion

We have documented for US states and selected countries with detailed census data that labor productivity and wages are both much lower in agriculture than in the rest of the economy. We then have established that: the productivity gaps are implausibly large compared to the wage gaps; there is a measurement problem in agriculture; the wage gaps are mostly accounted for by the fact that human capital is lower in agriculture. For US states we have established in addition that correcting for the measurement problem reduces the productivity gaps and makes them consistent with the wage gaps. This establishes that despite the large measured productivity gaps in US states, the US satisfies the assumptions that we typically make about the benchmark country in our models: after controlling for differences in human capital, wages are equalized across sectors and there is no room for beneficial policy interventions that reallocate labor from agriculture to non-agriculture.

Our results have two important implications for the development literature that seeks to account for large productivity gaps between non-agriculture and agriculture in developing countries. First, the fact that we find also for a our sample of countries that the wage gaps are mostly accounted for by large differences in human capital between non-agriculture and agriculture suggests that moving costs or barriers are not important quantitatively and that sectoral differences in human capital are the most promising explanation for productivity gaps between non-agriculture and agriculture. Second, the fact that even the BEA does succeed at precisely measuring agricultural productivity suggests that it is challenging to do this. Taken together with the evidence that measured productivity gaps are implausibly large also in our sample of countries, this suggests that there may well be a general measurement problem with agricultural productivity. This view is consistent with the fact that after a careful analysis of productivity gaps in a much larger cross section of countries than ours, Gollin et al. (2011) were unable to account for the entire productivity gaps between non-agriculture and agriculture. This suggests that more evidence is required to establish convincingly that the reported productivity gaps for the poorest countries are in fact well measured.

References

Adamopoulous, Tasso and Diego Restuccia, "The Size Distribution of Farms and International Productivity Differences," 2011. 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, "State Area Personal Income and Employment Methodology," Technical Report October 2009.

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.

____ **and Wilbur John Coleman II**, "How Regions Converge: Alternative Data Sources," June 1998. Manuscript.

____ **and** ____ , "The U.S. Structural Transformation and Regional Convergence: A Reinterpretation," *Journal of Political Economy*, 2001, *109* (3), 584–616.

Craig, Barbara J., Philip G. Pardey, and Johannes Roseboom, "International Productivity Patterns: Accounting for Input Quality, Infrastructure, and Research," *The American Journal of Agricultural Economics*, 1997, *79*, 1064–1076.

Donovan, Kevin, "Agricultural Risk, Intermediate Inputs, and Cross-Country Productivity Differences," 2012. Manuscript, Arizona State University.

Echevarria, Cristina, "A Three–factor Agricultural Production Function: the Case of Canada," *The International Economic Review*, 1998, *12*, 63–75.

Fishelson, Gideon, "Relative Shares of Labor and Capital in Agriculture: A Subarid Area Israel, 1952–1969," *Review of Economics and Statistics*, 1974, *56*, 348–352.

Fulginiti, Lilyan E. and Richard K. Perrin, "Prices and Productivity in Agriculture," *The Review of Economics and Statistics*, 1993, *75*, 471–482.

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.

Griliches, Zvi, "Research Expenditure, Education, and the Aggregate Agricultural Production Function," *The American Economic Review*, 1964, *54*, 961–974.

Hayami, Yujiro and Vernon W. Ruttan, "Agricultural Productivity Differences among Countries," *The American Economic Review*, 1970, *60*, 895–911.

Herrendorf, Berthold and Akos Valentinyi, "Which Sectors Make Poor Countries So Unproductive?," *Journal of the European Economic Association*, 2012, pp. 323–341.

—— **and Arilton Teixeira**, "Barriers to Entry and Development," *International Economic Review*, 2011, *52*, 573–602.

Internal Revenue Service, "Federal Tax Compliance Research: Individual Income Tax Gap Estimates for 1985, 1988, and 1992," Publication 1415 (Rev. 4-96), Washington, DC 1996.

——, "Reducing the Federal Tax Gap: A Report in Improving Voluntary Compliance," Technical Report, Washington, DC 2007.

Kislev, Yoav and Willis Peterson, "Economies of Scale in Agriculture: A Reeexamination of the Evidence," in John M. Antle and Daniel A. Sumner, eds., *The Economics of Agriculture, Volume 2: Papers in Honor of D. Gale Johnson*, University of Chicago Press, 1987, chapter 10, pp. 156–170.

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

Mundlak, Yair, "Economic Growth: Lessons from Two Centuries of American Agriculture," *Journal of Economic Literature*, December 2005, *43* (4), 989–1024.

——, **Donald F. Larson, and Rita Butzer**, "Determinants of Agricultural Growth in Indonesia, the Philippines, and Thailand," Policy Research Working Paper 2803 2002.

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.

Schultz, Theodore W., *Transforming Traditional Agriculture*, Yale University Press, 1964.

A Data Appendix: Sectoral Value Added

In this part of the Appendix, we explain in more 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).

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.[16] 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.[17]
- 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.
  - 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,

[16] The spreadsheets are available at `http://www.ers.usda.gov/Data/FarmIncome/Zip_filesXls.htm`.

[17] Note that NAICS uses animal production and USDA uses livestock production for the same industry.

some of the expenses of farms and some of the labor on farms are devoted to generating this revenue. Since we cannot isolate 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, in which case 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 counted as a factor payment and not as 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 and property taxes. To do so we take value added and add the line "net government transactions". Note that the government also provides indirect support for farmers through price supports and similar programs; the effects of these indirect supports are already counted in value added.

B Data Appendix: 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 Population Census

To calculate labor input from the US Population Census, we use the public–use census data made available through the IPUMS data service, Ruggles et al. (2010). We use the 5 percent sample for 1980, 1990, and 2000, which is the largest publicly available sample for 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 non-agriculture, we need to construct a crosswalk from workers' self-reported industry/occupation to our agriculture/nonagriculture classification. One potential complication is that the US 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 (e.g., "occ" and "ind" rather than "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., the animal and crop production industries. Table 10 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, non-agricultural workers are all workers with valid industry reports that do not work in an agricultural industry or in the military.

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 10. 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

Table 10: 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 11: 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.

with each Census and generally becomes more detailed over time. Table 11 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.2 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.[18] 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 1979, 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 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 total number of hours worked by this worker in all jobs, regardless of the sector. We cap total hours worked at 99 (consistent with earlier CPS procedures and a reasonable limit on the work week), and weight using the provided weight (pwsswgt). This gives us a measure of hours worked by primary workers in each sector, state, month, and year. We multiply this figure by 4.33 to generate monthly hours worked.

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

This measure of hours worked does not account for the time allocation of workers who have secondary jobs in a different sector than their primary job. In other words, these data are analogous to what we would find in the US Population Census, which does not distinguish between primary and secondary jobs. To account for secondary jobs in a different sector, we draw on data on primary and secondary jobs from the CPS from 1994 onward, and from the May supplements of the 1979, 1980, 1989, and 1991 CPS. We use this information in the following way:

1. For the months where the data are available, we calculate the fraction of the total hours worked that is devoted to farming by workers whose first jobs are in farming and non-farming. We aggregate this information to the state–month–year level.

2. We use a regression with state and month fixed effects as well as a linear time trend to predict the time allocation of the workers with a primary job in farming and non-farming for years in which the data are not available.

3. We combine our information on the hours worked by those with primary jobs in farming and non-farming with our predicted time allocations of hours to calculate labor in farming and non-farming.[19]

We measure wages per hour using CPS data to be consistent with our other work. We perform wage regressions using the wage data from the NBER Matched Outgoing Rotation Group (MORG) during 1980–1993 and from the outgoing rotation groups in the CPS Basic Monthly Data during 1994–2009. We again use the outgoing rotation group weights and we multiply top–coded wages by 1.4 as suggested by the CPS. As is standard, we run the wage regressions for a selected sample of workers who meet the following criteria:

- They work for wages and salaries (the reason for this restriction is that the reported wages of self-employed or unpaid workers are considered unreliable, and are usually not even collected).

- They have a valid hourly wage, that is, they either 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 working at least 30 hours per week.

[19] Note that this means that we use the predicted hours also when we actually have hours in the CPS. We do this for logical consistency and to smooth the data in states with small farming samples.

- 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 are 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.

One final issue to address is that monthly CPS wage data do not include all the compensation that labor receives. First, they do not include irregular compensation, such as bonuses. This type of income is better captured in retrospective questions, such as the question in the March CPS supplement on total labor income earned in the last year. We construct the ratio of retrospective March CPS income to the sum of all monthly income in the prior year, and use this ratio as a correction for irregular compensation. The second component of labor compensation missing is benefits. Fortunately, NIPA includes information on wages and on total compensation (wages plus benefits) by state, year, and industry. We correct CPS wages by the ratio of NIPA total compensation to NIPA wages to adjust for benefits. This correction is generally slightly larger for the non-agricultural sector, indicating that benefits are more generous there. All our figures on wages in the paper include both of these corrections so that they represent total labor compensation.

C Data Appendix: Proprietors' Income

C.1 Census of Agriculture

We calculate revenues minus costs for farm proprietors using the 2002 and 2007 Censuses of Agriculture. The Census of Agriculture distinguishes between family or individual farms, partnerships, corporations, and other farms (which includes cooperatives, trusts, institutional farms, and other unusual arrangements). We calculate farm proprietor income as farm income less farm expenses and factor payments for family or individual farms and partnerships. In theory, the result should be identical to farm income from the March CPS, which measures the income from owning and operating one's own farm, unless one is incorporated. In practice, it is quite close.

We construct income produced on farms as follows:

- The value of total sales is farm income. This includes the sales of crops and livestock of many types.
- The value of government payments is farm income. This category captures subsidies from the federal government.
- The revenue from several services is farm income:
  - 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.
  - The sale of forest products is farm income, for the same reasons as it is counted in the USDA.
  - Patronage dividends and refunds from cooperatives are farm income, which occur when farmers paid for inputs through cooperatives but price realizations come out lower than expected.
  - Agri-tourism payments, insurance payments, and payments from state and local government programs are farm income.
  - Other farm-related income is farm income.
- We do not count one other type of farm income, namely gross cash rent or share payments, which should be reported as a factor payment elsewhere.

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

- The purchases of products produced in the farm sector (including seeds, livestock, and feed) are farm expenses.
- The purchases of manufactured inputs (including fertilizer, lime, chemicals, gasoline and other fuels, and utilities) are farm expenses.
- We count repair and maintenance of capital items, machine hire and custom work, and miscellaneous expenses as farm expenses.
- We do not count as farm expenses several items that are factor payments, such as rental payments for land and machinery; payments to landlords; interest payments; and payments to workers.

Finally, we take the measure of property taxes directly from the line "property taxes paid". Our measure of revenues minus costs of proprietors then is farm income minus farm expenses, farm factor payments, and property taxes paid.

C.2 March supplement of the CPS

The March supplement to the Current Population Survey asks about the pre-tax income of farmers who own and operate their own farm or are self-employed on their own farm. It does not contain information on the income of those who work as salary or wage employees, including those who work on their own incorporated farm; this information is collected elsewhere. Moreover, it does not contain information on the earnings of owner non-operators, such as those with shares in a farm corporation. Thus the pre-tax income of farmer from the March supplements corresponds closely to the figure computed from the Census of Agriculture above.

We download the relevant data from the IPUMS data service, Ruggles et al. (2010). The relevant variable is incfarm. We add this income to the national total using the provided weights (wtsupp).

D Data Appendix: Consumer Price level

To construct the consumer price level in agriculture and non-agriculture in each state, we use the price parities that Aten (2006), Aten and D'Souza (2008), and Aten (2008) estimated for the year 2006 for 363 metropolitan areas and for the rural area of each 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.