March 28, 2017

Heterogeneity, Measurement Error, and Misallocation: Evidence from African Agriculture

Douglas Gollin
Oxford University

Christopher Udry*
Yale University

* Authors' addresses: Oxford Department of International Development, Queen Elizabeth House, 3 Mansfield Road, Oxford OX1 3TB, United Kingdom; Department of Economics, Yale University, 27 Hillhouse Ave., PO Box 208269, New Haven CT 06520-8269.
E-mail addresses: *douglas.gollin@qeh.ox.ac.uk* and *udry@yale.edu*. We are grateful to Philipp Wollburg for excellent research assistance. We have received helpful comments from seminar participants at Princeton, Heidelberg, Yale and CEMFI-Madrid. We are grateful for comments from Andrew Foster.

1. Introduction

How important is misallocation in explaining the income differences across countries? A recent literature in development and growth economics has focused on misallocation across sectors, firms, and plants.[1] This literature has found evidence that the dispersion of productivity across production units seems to be consistently higher in poor countries than in rich ones. Such productivity differences have the potential to account for a large fraction of the cross-country income differences. In an aggregate sense, misallocation across sectors or firms reduces aggregate total factor productivity (TFP).

A challenge in this literature is to distinguish misallocation from other sources of dispersion in productivity, such as technology shocks, measurement error, and adjustment costs of various kinds. This paper seeks to disentangle these different sources of dispersion in an environment where cross-firm differences in productivity are very large, aggregate productivity is low, and market failures undoubtedly contribute to cross-firm frictions in the allocation of resources. Specifically, we take advantage of extraordinarily rich data from farms in three countries in Africa, for which we have detailed panel observations on the same firms producing identical outputs on different plots in the same time period. Since farmers face no market imperfections in allocating resources across their own plots within a growing season, we can use these plot-level data to identify misallocation more precisely. Our strategy allows us to disentangle the productivity dispersion that arises from misallocation from that stemming from measurement error or heterogeneity in technology and inputs (including production shocks).

The agricultural sector provides a valuable window through which to study firm-level misallocation. Most firm surveys have relatively few observations on different plants or establishments operated by the same firm, and this makes it difficult to disentangle firm management from any unobservable characteristics of the plant or factory. Another advantage we have relative to firm surveys is that our producers are producing highly homogeneous products, with little market power. Consequently, we can compare the output of different firms (farms) without worrying about markups and pricing strategies.

[1] See, for example, Hsieh and Klenow 2009; Restuccia and Rogerson 2008, 2013, McMillan et al. 2014.

Understanding the extent of misallocation in agriculture is also of interest because of evidence that low agricultural productivity can explain – at least in a mechanical sense – a large fraction of the cross-country dispersion of output per worker (Caselli 2005, Restuccia et al. 2008). A cluster of recent papers has suggested that there may be very large dispersion in productivity at the level of farms and farmers, potentially indicative of misallocation at this micro level.[2] These papers point out that in poor economies, very large fractions of the workforce are employed in agriculture, in contrast to rich countries, where very few people earn a living from farming. In economies where two-thirds of the people are farmers, it is reasonable to ask whether they are all good at farming – and whether market failures of various kinds may induce too many low-skill farmers to remain in agriculture.

Restuccia and Santaeulalia-Llopis (2017), in particular, have raised the intriguing possibility that much of Africa's productivity deficit might be attributable to misallocation within the agricultural sector. In particular, they find suggestive evidence, based on data from Malawi, that too much farmland is managed by low-skill farmers. If true, this finding might offer an explanation for sub-Saharan Africa's low productivity in agriculture. Indeed, it might by extension help explain the region's low levels of income per capita. The finding also suggests a relatively straightforward solution – albeit one with great political complexity – namely, the liberalization of land and input markets, so that the best farmers can eventually buy out those farmers who lack the skill to farm productively.

The misallocation hypothesis for African agriculture is particularly plausible because of abundant evidence that the continent's agricultural markets work poorly – for land, rural labor, intermediate goods, and output. Much land lacks formal title, and rural labor markets are often poorly integrated. Empirical tests consistently reject the hypothesis that markets are complete.[3]

---

[2] See, for example, Adamopoulos and Restuccia 2014, Adamopoulos and Restuccia 2015, Adamopoulos et al. 2017, Bento and Restuccia 2016, Restuccia and Santaeulalia-Llopis 2017.

[3] See, for example, Dillon and Barrett (2017) for a set of African countries; Karlan et al. (2014) for Ghana, Udry (1994a) for Burkina Faso, and Udry (1996) for Burkina Faso and Kenya. Similar findings are common for other parts of the developing world as well; see LaFave and Thomas (2016) for Indonesia; Kaur (2016) and Jayachandran (2006) for India.

At the same time, market failure need not lead to misallocation. Development economists have repeatedly and convincingly documented the existence and effectiveness of rural institutions that can stand in for complete markets, with at least limited effectiveness. Informal credit markets appear to substitute imperfectly for both formal credit markets and formal insurance markets.[4] This literature has argued that informal institutions can often succeed in avoiding gross inefficiencies – perhaps as the result of some evolutionary pressures that shape these institutions over time. From this perspective, the persistence of very costly land misallocation across farmers would pose a puzzle.

Our paper addresses the measurement of misallocation using panel data from three countries (Ghana, Tanzania, and Uganda) for which we can observe production in great detail. In these data, we can observe the inputs and outputs for specific crops cultivated by individual farmers – not simply households -- on specific plots of land. The data are similar to those used by Restuccia and Santaeulalia-Llopis (2017), although we exploit the panel dimension of these data sets rather than the cross-section. For each of our three countries, we can observe many of the same individual farmers in at least three periods.

The rich detail of the data allows us to disentangle misallocation from three other important sources of variation in measured productivity at the farm level. The first of these is simply the stochastic nature of agricultural production. Farmers face a large number of idiosyncratic shocks to production that are not well observed in the data, related to weather, pests, crop diseases, and so on. A second source of variation in productivity is measurement error; in spite of the high quality of the data that we work with, reporting is imperfect and measurement is imprecise.[5] Finally, the third source of variation in productivity is heterogeneity in unobserved land quality.[6] All will give rise to dispersion in measured total factor productivity (TFP) at the farm level, as well as to dispersion in input intensity. Because of this, any estimates of the potential gains from reallocation need to account carefully for mismeasurement and heterogeneity.

[4] Early papers in this literature included Townsend (1994) and Udry (1994b).
[5] See, for example, De Nicola and Giné 2014, Deininger et al. 2012, and Beegle et al. 2012b; although Beegle et al. 2012a offer a more positive view.
[6] The problem of unobserved land quality was recognized by Benjamin 1995 and Udry 1996. More recent surveys often collect quite detailed data on soil quality, but the dimensionality of soil quality measurement can be overwhelming; see, for example, Tittonell et al. 2008.

In this paper, we propose a theoretical framework that models the processes by which farmers select plots, allocate inputs to individual plots, and subsequently realize output. Our theoretical framework explicitly recognizes the stochastic nature of agricultural production and the sequencing of farm decision-making. We then show how this model can help distinguish empirically between misallocation, mismeasurement, and heterogeneity, given plot-level data.

Drawing on the model, we assess the relative importance of the three different explanations. Our results suggest that idiosyncratic shocks, measurement error, and heterogeneity in land quality are important sources of dispersion in productivity across farms. We find that when these are taken into account, the overall importance of misallocation drops substantially. We estimate that reallocation of land to the most productive farmers has the potential to increase aggregate output by about 15 percent for Ghana and about 50% for Uganda. (For Tanzania, we carry out a preliminary analysis of dispersion, but we do not yet have estimates for reallocation.) These are large gains, to be sure, but they are sharply lower than estimates of reallocation effects that do not account for these confounding sources of dispersion. Although a gain of 15-50 percent would clearly be quite significant, our results suggest that efficient reallocation of land and other agricultural inputs would not dramatically close the income gaps between African countries and the world's rich economies.

Although our work focuses on agriculture, many of the same issues clearly matter for the broader literature on the importance of misallocation across firms in the developing world. Much of the macro literature on misallocation has abstracted entirely from measurement error and heterogeneity. Our results show that estimates of the gains from reallocation are highly sensitive to assumptions about these other sources of dispersion in measured productivity. Put simply, if most of the cross-firm dispersion in productivity arises from manager characteristics, the gains from reallocation will be large. But if most of this dispersion arises from (unobservable) shocks and mismeasurement, then reallocating assets to different firms and managers will have little impact.

An important caveat of our work is that we consider only the effects of static misallocation. Implicitly, this holds constant the existing institutions and

technologies. With improved technologies and different institutions, one might expect that the efficient allocation of land and inputs across farms and farmers would look very different. For instance, with different market structures and institutions, farmers in our three countries might find it worthwhile to mechanize more fully and to use tractors for land preparation and other farming activities. Given a shift from human power to mechanical power, the efficient operational size of a farm might change quite dramatically, and labor might be replaced by capital, with farm size increasing as it has in Europe and North America. Our analysis does not consider this hypothetical case. Neither do we ask whether technology adoption would take place more rapidly if farms were consolidated. In this sense, our results are not necessarily inconsistent with those of Adamopoulos and Restuccia (2014), who ask how agricultural production would change if all countries had the same size distribution of farms that is observed in the United States.

The remainder of this paper proceeds as follows. Section 2 provides some descriptive background and reviews related literature. We show how our paper connects to a number of strands in both the micro and macro literature. Section 3 presents some descriptive features of the data. In particular, this section shows that for our three countries, the dispersion of productivity across farms remains largely unchanged as we zoom in from the national level to increasingly disaggregated geographic levels. We show that dispersion in productivity at the district and village levels are almost as large as that at the national level. Even within farms, the dispersion on different plots cultivated by the same farmer is large relative to the dispersion in productivity across all farms at the national level. We interpret this as implying that farmer characteristics do not accounting for the bulk of the dispersion in productivity that we observe in the national data. Moreover, we find that there are important patterns of productivity across plots within farms. In Section 4, we draw on these patterns of productivity variation within farms to construct a theoretical framework that models the ways in which farmers choose their plots, select the crops (or crop combinations) that they cultivate on each plot, apply inputs, and realize output. In Section 5, we use this model to motivate a structural estimation of agricultural production functions for our three countries. Finally, in Section 6, we use the estimated production functions to model the impact of a reallocation of land to the most productive farmers. Section 7 discusses these results, and Section 8 concludes.

2. Background and literature review

Across sub-Saharan Africa, over 60 percent of the population lives in rural areas, and agriculture remains the dominant source of employment in most countries of the region (World Bank, World Development Indicators; henceforth WDI). Measured productivity levels are extremely low. Value added per worker in African agriculture appears to be less than half the level attained in other sectors, even after adjusting for differences in input quantity and quality (Gollin et al., 2014). In a proximate sense, these two facts imply an unpleasant agricultural arithmetic for African income levels: if many people earn their living from agriculture, and if agricultural incomes are low, then aggregate incomes will be correspondingly low.

The disparities in average labor productivity across sectors do not necessarily imply misallocation, however. Average productivity is not the same as marginal productivity, so sectoral differences in productivity could arise efficiently from sectoral differences in capital intensity, to give one example. Average labor productivity could also differ across sectors due to unobserved differences in worker skills. For instance, higher-skill individuals might tend to leave agriculture, so that average productivity would differ across the two sectors -- but there might be no difference for a worker of a particular skill level.

Nevertheless, there are many reasons to consider seriously the possibility that misallocation could be an important factor in explaining sectoral differences in productivity. For a start, the sheer number of people working in agriculture suggests the possibility of misallocation. In rich countries, only one or two percent of the workforce will be engaged in farming; in sub-Saharan Africa, nearly two-thirds of the workforce consists of farmers. Presumably not all the people working in African agriculture are particularly gifted as farmers. Some will surely be better than others. But for a variety of reasons, many people born in rural areas find it difficult to leave, and rural institutions in much of sub-Saharan Africa are designed to share farmland and other resources with those who remain.

This view of misallocation has motivated a series of recent papers that have explored the possibility that there are too many small farms in the developing world, with too many of these farms operated by poorly skilled farmers. This view is at the heart of work by Adamopoulos and Restuccia (2014, 2015) and Restuccia and Santaeulalia-Llopis (2017), among others. These papers explore the hypothesis that distortions in farm size may account for a large fraction of cross-country differences in agricultural productivity. Similar issues are explored in Adamopoulos et al. (2015), Chen (2016), Emran and Shilpe (2015), Gottlieb and Grobovsek (2016).

This literature builds on a broader literature in growth economics that has emphasized the importance of misallocation across firms and plants as a potential source of cross-country productivity differences or macro fluctuations in rich countries; e.g., Syverson (2004, 2011), Petrin et al. (2011), Petrin et al. (2013). An important branch of this literature has viewed misallocation as a plausible -- and indeed likely – explanation for low aggregate productivity in widely varying contexts, including developing countries; e.g., Banerjee and Moll (2010), Bento and Restuccia (2015), Garcia-Santana and Pijoan-Mas (2014), Hopenhayn (2014), Hsieh and Klenow (2009), Kalemli-Ozcan and Sorensen (2010), Midrigan and Xu (2014), Restuccia and Rogerson (2008, 2013). A recurring theme of this set of papers is that the misallocation of productive resources into low productivity firms can lead to low aggregate productivity. Empirical analysis generally supports the idea that poor countries have many firms with low measured TFP. The reasons for the persistence of these low productivity firms are not always clear, but a sufficient explanation would be frictions or policies in poor countries that induce distortions to the efficient size distribution of firms.

A challenge in this literature is the measurement of productivity at the level of individual firms. Typically, the data used for these analyses come from firm surveys that may vary in quality and in coverage. To calculate measures of productivity for the individual firm requires a series of strong assumptions about the firm-level production function and about the quality of data. In particular, methods used widely in the macro literature on misallocation have been criticized on methodological grounds; e.g., by Asker et al. (2014), Foster et al. (2016), and Haltiwanger (2016). Our approach addresses some of the concerns raised by these critiques. In particular, our approach recognizes that idiosyncratic shocks (such as

weather shocks) and adjustment costs (such as rigidities arising from early-season planting decisions) may give rise to dispersion in productivity that is consistent with efficient allocations rather than misallocation.

Our paper also connects with a long strand of micro development literature that has examined some of the same questions around allocative efficiency that have been taken up in the recent macro misallocation literature. The literature on efficiency within and across farms in developing countries dates back to a large literature on the rationality of farms in developing countries, starting perhaps with Chayanov's work on peasant economies (republished in English 1966) and taken up again in the work of Schultz (1964). A large literature from the past half century has tried to understand efficiency in the context of agricultural household models. This literature has explored the ways in which agricultural households facing incomplete markets may make choices that are efficient subject to a variety of constraints. Our theoretical framework and analytic approach are consistent with this literature.

3. Data and Settings

Our paper draws on three different data sets. Two are nationally representative data sets, for Tanzania and Uganda, collected by government statistical agencies in collaboration with the World Bank's program on Living Standards Measurement Surveys – Integrated Surveys of Agriculture (LSMS-ISA). The first of these is the Uganda National Panel Survey (UNPS), which has followed about 3,200 households that were interviewed in 2009-2010, 2010-2011, 2011-2012, and 2013-14. The second is the Tanzania National Panel Survey (TZNPS), which has followed about 3,300 households that were interviewed in 2008-2009, 2010-2011, and 2012-2013. The third is a survey that is representative at the sub-national level for northern Ghana, carried out by Innovations for Poverty Action (IPA) as part of its project on Examining Underinvestment in Agriculture, which followed 1,358 households in an unbalanced panel over three years. All three surveys collected data on the individual plots cultivated by particular farmers within households, and detailed information were collected by plot on inputs used and output harvested. Depending on the survey, some or all plots were measured

by GPS, and data were collected using state-of-the-art survey techniques. The data are freely available online and all data and documentation are available for open access.[7]

Data for all surveys include detailed descriptors of both the households and the farms. For households, data are available on household composition and the age, education, and health characteristics of each household member; the relationship of each member to the household head; and the allocation of each person's time to household production and market labor, among many other variables. For the farm, data were collected at the plot level on crops cultivated, soil characteristics, toposequence, location, soil quality (including measures of erosion and tree cover), land rights, and a variety of observed shocks, including rainfall. We note that in addition to all these aspects of the data, we have additional information for the Ghana sample about the responses of some households to experimental variation in farmer-level budget constraints, based on a randomized controlled trial (RCT) that was conducted as part of the EUI project. This variation is used in our identification strategy as described in Section 7.

An important feature of our data – and one that helps us significantly in terms of our identification strategy – is that we have many instances in each country in which we observe the same farmer cultivating the same crop on multiple plots within the same year. For instance, we may observe a single farmer growing maize on each of two or three distinct plots in the same growing season.[8] This is not particularly surprising; in many African production environments, farmers may farm non-contiguous plots because of complex patterns of inheritance and land rights. Even when the plots are contiguous, farmers may plant different plots with the same crop but at different dates or with different varieties, due to the

---

[7] For information on the LSMS-ISA project and links to the data, see: http://go.worldbank.org/BCLXW38HY0.

[8] For convenience, we speak of "a farmer" as an individual. But our data sets actually provide quite rich data that distinguishes the person who owns the land from the person who manages the plot and the person who keeps most of the revenue from the plot. We focus here on the person who manages the plot. An added level of complexity is that the data often allow for up to two household members to be designated as the manager of the plot. We use the term "farmer" to refer to distinct individuals or pairs of household members. When we speak of a farmer cultivating the same crop on different plots, it could thus be a husband and wife (or father and son, or two brothers,etc.) operating as a pair.

micro characteristics of the plots or as an effort to diversify against shocks that might occur at different points in the season.

Our three countries differ to some degree in the types of production systems that we observe. Some crops are common across the three countries (e.g., maize), while others (*matoke*, a kind of cooking banana) are of importance only in a single country (in this case, Uganda). Input use differs across countries, too: in Uganda and Tanzania, labor is essentially the only input used by most farmers. By contrast, farmers in the Ghana sample use substantial amounts of purchased fertilizer and pesticides.

3.1 Descriptive statistics

Tables 1a and 1b show key descriptive statistics for our three data sets. As Table 1a shows, within some households, there are multiple farmers; the individual manager is documented for each plot. For instance, within the Uganda data, the 2,592 farm households correspond to 3,671 distinct farmers. We observe these farmers over six seasons, and we end up with nearly 40,000 plot-season observations. For Tanzania, we have almost 17,000 plot-season observations, and for Ghana, we have over 8,000. Individual plots are quite small, ranging from 0.20 ha in Uganda to 1.21 ha in Ghana. The majority of farmers cultivate multiple plots within each season. Thus, for Uganda, the median number of plots per farmer-season is 4; for Ghana, it is 3; and for Tanzania, the median farmer cultivates 2 plots per season. Not all of these plots are cultivated with the same crops; the median number of plots that a given farmer cultivates with a given crop in a given season is one.

Our samples are geographically quite dispersed for Uganda and Tanzania, covering over 600 villages across 183 districts in Uganda and 184 villages across 140 districts in Tanzania. The Ghanaian sample is narrower, drawn from 75 villages in just three districts.

Table 1b shows yields (output per hectare) for each of the data sets. These are given in value terms because of the prevalence of multiple cropping (i.e., several crops being cultivated at the same time on a given piece of land). Multiple cropping makes it difficult (or irrelevant) to measure yield in physical quantities.

Instead, we report value per hectare, with the physical quantities of different crops priced using median values reported by all farmers in a community.

It is immediately apparent from the yield data that yields are wildly skewed. The mean yield for is typically around twice the median, and the large standard deviations are indicative of very long right-hand tails of the distributions. This is true even after the data have been winsorized at the 0.01 level. Because there are biophysical constraints on maximum yield, we look skeptically at some of the very high reported values of yield in these data, and we view this as *prima facie* evidence that measurement error is likely to be an important feature of the data.

The data on input intensity are somewhat less skewed, especially for Ghana and Uganda. Only in Ghana is there significant use of tractor ploughing for land preparation, with tractors used on 70 percent of plots. We define the land input in Ghana as prepared fields, whether ploughed by hand or by tractor and the labor input is defined as post-land preparation labor. Median days of labor per plot are not very different across the three countries, however; with 33 days per plot in Uganda, 39 in Ghana, and 43 in Tanzania.

4. Heterogeneity, allocative efficiency, and variation in the intensity of cultivation

In this section, we document the dispersion in measured productivity across farms and plots, and we explore patterns that are evident in the data. It is useful to begin with a simple benchmark model of efficient static allocation.

4.1 Efficient static allocation

Consider a population of farmers indexed by h, with each farmer cultivating a set of plots indexed by i. Production (Y_{hi}) on each plot depends on inputs of land (L_{hi}) and labor (X_{hi}), as well as known plot productivity (ω_{hi}) and unexpected shocks to output (ϵ_{hi}) according to the concave production function

$$Y_{hi} = f(L_{hi}, X_{hi}; \omega_{hi}, \epsilon_{hi}).$$

The distinction between ω and ϵ is that the farmer knows ω_{hi} at the time she chooses (L_{hi}, X_{hi}), but at that time she only knows the distribution of ϵ.

If the allocation of resources were efficient, then there would exist a set of common prices such that, with respect to those prices, all farmers are maximizing profits on each of their plots. In such a world, if $f(L, X; \omega, \epsilon)$ is homothetic in (L, X), then perfectly-measured factor ratios would be identical across all plots. If in addition, there were no risk ($\epsilon_{hi} = 0, \forall h, i$), then perfectly measured output-factor ratios would also be identical across all plots.

Needless to say, this description does not characterize the world particularly well, and in particular our data show marked deviation from this benchmark. There is wide dispersion in factor ratios across plots (e.g., labor per unit land, fertilizer per unit land) as well as in realized output per unit land. This dispersion is large and ubiquitous, and it remains even after controlling for a variety of observable plot characteristics and observable shocks.

4.2 Dispersion of yield and factor intensity

Consider first Figure 1. Each subgraph of this Figure shows a set of Epanechnikov kernel density estimates of the density of the deviation of log output per hectare from its sample mean. Each subgraph corresponds to one of the three countries in our data. The different lines on the Figure correspond to dispersions calculated with differing controls.

Figure 2 presents similar estimates of the density of the deviation of log labor per hectare in each country. This is a measure of input intensity, which is a useful alternative to the measure of realized yield. One might imagine, for instance, that yield is a noisier measure, given that output realizations necessarily embody all the shocks that have occurred during the growing season. By contrast, much of the labor applied to each plot is realized before harvest and hence should not reflect all of the shocks that might alter yield.

Consider first the solid black lines in Figure 1 and Figure 2. These are the raw dispersions across plots. Figure 1 shows that the variance of log output per hectare ranges from 1.08 in Ghana to 1.98 in Uganda. The variance of log labor input per hectare ranges from 0.74 in Ghana to 1.08 in Tanzania. The lower variances of these ratios in Ghana reflects, at least in part, the greater homogeneity of the farming systems in the Northern Region of Ghana compared to those of Uganda or Tanzania as a whole. It is noteworthy that the variance of log labor input is

quite high; yield dispersion is not coming entirely from shocks affecting final harvest.

All three data sets contain rich information on observable components of both ω and ϵ. The first task, therefore, is to account for observable heterogeneity.

Land characteristics such as slope, soil type, and location contribute to ω and are measured in each of our data sets. Characteristics of the farmer such as gender, education, and experience are also components of plot productivity that we observe. Agriculture in each of these three settings is almost exclusively rainfed. Realizations of the total level of rainfall and its distribution over the season contribute to both known plot productivity and to unanticipated shocks to output, depending upon the timing of the realization. We condition on measures of these shocks, and their interaction with land characteristics as well. If these observed characteristics fully account for the variation in ω and $\epsilon,$ then output per hectare and labor per hectare will not vary across plots in an efficient allocation, once we control for observables.

Tables 2(a)-(c) reports a set of regressions for each of the three countries, with output per hectare as the dependent variable in all regressions. Observations are for individual plots in specific years/seasons. In each of these tables, the first column shows selected coefficients from a regression of output per hectare on cultivated area and the large set of observable land characteristics and exogenous shocks that are available in these datasets. The estimated density of the residuals from these regressions is illustrated as the red line in each of the subgraphs of Figure 1. The plot characteristics and shock variables are highly jointly significant in each regression, and the estimated variance of the residuals is significantly smaller than the variance of the raw data in each case. This tells us that the observable plot characteristics are indeed explaining part of the dispersion in yield. Nevertheless, as is apparent from Figure 1, including these observable plot characteristics does not alter the overall pattern of dispersion in productivity.

The first column of Tables 3(a)-(c) reports the same subset of coefficients of the parallel regression of labor input per hectare on the observable land characteristics and exogenous shocks. The estimated density of the residuals from these regressions is the second line in each of the subgraphs of Figure 2. Again, the set

of observable characteristics is highly jointly significant in each regression, and the estimated variance of the residuals is significantly smaller that of the raw data. The variation in observable characteristics, including shocks, is an important determinant of the variation in both output and labor input per hectare across plots in each of these samples. But again, these observables do not generate much difference in the pattern of residuals as shown in Figure 2.

Table 4 and Figure 3 repeat this exercise for the intensity of use of post-planting inputs (largely fertilizer and pesticides) in Ghana, the only sample in which purchased inputs are used to an appreciable degree.[9]

In the analysis thus far, the within-country data pool observations across farming systems and over multiple growing seasons. Differences in technology across farming systems and crops and variation over time in the shadow costs of factors of production or the shadow value of output could generate variation in output or labor per hectare even in an efficient allocation with homothetic production functions. Therefore, Column (2) in each of the Tables 2(a)-(c) and 3(a)-(c) reports coefficients from regressions of log output per hectare and log labor per hectare on the same set of plot characteristics with year-season-region-crop fixed effects.[10] Estimates of the density of the residuals from these regressions are the third lines graphed in each of the subgraphs of Figures 1 and 2.

Qualitatively speaking, these tables provide evidence that observable characteristics of plots and shocks have a statistically meaningful effect on input intensity and yield. However, we note that *quantitatively* speaking, these observables do not account for a very large fraction of the total dispersion.

One way to see this is to note that the magnitude of the remaining variation is large: the variance of the log residual ranges from 0.71 in Uganda and 0.75 in Ghana to 0.82 in Tanzania. In comparison, the variance of the deviation of log

---

[9] In what follows, we sometimes refer to these purchased non-agricultural inputs as “capital,” to distinguish them from labor, which is the principal non-purchased input. Chemicals are not strictly capital in the sense that they cannot be accumulated as stocks over time. But they must be paid for out of cash, in contrast to most of the other inputs and factors of production.

[10] The Ghana regressions reported in Table 2(a) and Table 3(a) do not include a region fixed effect, because all observations are in the Northern Region, which is a relatively uniform production environment.

output per hectare for farms in the United States is 0.05 for corn in the Corn Belt and 0.14 for wheat in the Northern Plains (Claassen and Just 2011).[11] The variance of log labor per hectare also remains substantial: it is 0.43 in Tanzania, 0.32 in Ghana and 0.29 in Uganda.

It is apparent that substantial variance in output per hectare and labor input per hectare remains after we account for a rich set of observable characteristics of land, including detailed measures of rainfall variation. The variance remains large even when we add crop-season-region fixed effects (Column 2) and indeed when we add fixed effects at a more narrowly defined geographic or administrative level. For example, in all the tables, Column (3) adds community or village fixed effects. This remaining variation is sometimes characterized as reflecting the effects of factor and output market distortions that prevent the efficient match of factor inputs to dispersion in total factor productivity (Hsieh and Klenow 2009; Adamopoulos et al., 2017; Restuccia and Santaeulalia-Llopis 2017). For this reason, we will refer to the estimated residuals from the regressions reported in columns (2) of Tables 2(a)-(c) and 3(a)-(c) as our baseline measures of dispersion in productivity across plots.[12]

However, this baseline dispersion might also be a consequence of *unobserved* characteristics of land; it might also reflect unobserved dimensions of risk, or measurement error in output or factor inputs, even if factors of production are allocated with full efficiency. In order to draw useful conclusions regarding the extent of factor misallocation and its implications for aggregate output loss, it is vital to disentangle these sources of variation. To do so, we rely on an assumption that *within a farm*, the allocation of resources across plots is efficient.

[11] Claassen and Just (2011) report that for more than 500,000 observations in their US data, the 95th percentile corn yield is 190% higher than the 5th percentile, which they view as "quite wide" (p. 148). By contrast, we find 90-10 ratios for Ghana of 740%, for Uganda of 2630%, and for Tanzania of 2400%. This reinforces our perception that the dispersion of yield across plots is very high.

[12] An alternative baseline could be provided by examining the residuals from a similar regression with village-crop-year fixed effects. This would absorb the effects of unobserved village-level shocks which might otherwise be misinterpreted as misallocation, but it would also absorb any real misallocation of resources across villages. As can be seen in Figures 2-3, the estimated dispersion of the residuals from these two specifications is similar.

A farm is defined as the set of plots cultivated under the management of a single farmer in a single season. Any reallocation of factors across plots within a farm requires no market intermediation or other exchange, only rational decision-making by the farmer. While we acknowledge that there may in fact be behavioural limits on the rationality of input decisions by farmers, we abstract away from these sources of efficiency loss for this paper and maintain the Schultzian "poor but efficient" assumption. This assumption does not imply that farmers are equally productive or knowledgeable. One farmer may have superior technical knowledge to another; this difference would be reflected in higher total factor productivity of the first.

If the allocation of factors across plots within a farm (during a single season) is efficient, then the dispersion of factor- and output-factor ratios across plots within a farm is generated by (a) imperfect measurement of factor inputs; (b) imperfect measurement of output; (c) varying realizations of risk; or (d) violation of the assumption that the production function is concave and homothetic. We will not consider (d) further.[13]

The final two columns of Tables 2(a)-(c) and 3(a)-(c) show coefficients from regressions of log output per hectare and log labor per hectare with the same set of plot characteristics and within-farm fixed effects. To be precise, Column (4) in each table reports the regressions with crop-season-household fixed effects, and Column (5) is based on crop-season-farmer fixed effects, where we are now looking at variation across plots farmed by the same individual or pairs of individuals from within the households. The residuals from these regressions are again shown in each of the subgraphs of Figures 1 and 2.[14]

In each country, when we consider the yield regressions of Table 2, approximately one-third of the baseline dispersion from the specification reported

[13] It is of course possible that some farmers are systematically worse than others at allocating efficiently across plots. But it is not obvious that this should have a strong correlation with productivity *levels*. A bad farmer is arguably one who realizes equally poor yields across all plots, based on allocating inputs with the same (improper) intensity across all plots.

[14] The fifth, penultimate column reports the results of similar regressions with household-crop-season fixed effects. There is evidence of systematic differences in yield and labor intensity on the plots of men and women farmers within the same household in Ghana, but not in Uganda or Tanzania. Even in Ghana, the magnitude of the dispersion generated by this difference is very small relative to other sources of variation, as found by Udry (1996).

in Column (2) remains after we focus attention on variation in output per hectare across plots *within a farm* -- and even across the plots farmed by an individual farmer. When we look instead at the labor intensity regressions of Table 3, one-fifth (in Ghana) to one-third (in Uganda) of the baseline dispersion remains after we restrict attention to variation within farms.

Given our assumption of efficient within-farm allocation, we conclude that this residual variation is evidence for significant measurement error in factors of production or output. Alternatively, it could reflect unobserved shocks to output that do not affect the marginal product of factor inputs or which occur following the application of inputs to different plots within a farm. If the variance of these errors of measurements or of shocks to output is at least as large across farmers as it is across plots of a given farmer, then attributing the residuals of the equations estimated in columns (2) of Tables 2 and 3 to misallocation would result in an overstatement of the gains to reallocation.

An estimate of the gains to efficient reallocation of factors across plots requires knowledge of the production function and of the magnitude of measurement error in factor inputs, which we address in Section 5. However, we can generate a rough calculation of the degree of overstatement of the potential gains to reallocation generated by the baseline estimates' overstatement of the dispersion of log productivity. As we show in Appendix QQ, the gain to an efficient reallocation of factors across plots is increasing in the exponential of the variance of the residual log output per hectare, conditional on truly measured land inputs and output, and on all shocks to output.[15]

In all three countries, the residual variance of log output per hectare within farm–crop–year groups is approximately 0.25. If this variance is generated by measurement errors in output or land, or by late season output risk, and if those factors have a similar variance across plots of other farmers, then the correctly measured residual log variance is approximately 0.25 units lower than the baseline estimate presented in Column (2) of Tables 2 and 3. The impact of this

[15] If the production function is Cobb-Douglas with coefficients α_L and α_X, and ϵ is distributed normally, then the gain from reallocation is directly proportional to $\exp\left(\frac{\sigma_\omega^2}{2(1-\alpha_X-\alpha_L)^2}\right)$, as shown in the appendix.

adjustment on the gain to reallocation depends upon the concavity of the production function. If the production function is homogeneous of degree 0.6, as assumed for example by Restuccia and Santaeulalia-Llopis (2017), then the overstatement of the variance of the residual by 0.25 generates a 190 percent overstatement of the efficient level of output. If the production function is homogeneous of degree 0.85, as we estimate below, the overstatement of the variance of the residual results in a 550% overstatement of the efficient level of output.

4.3 Plot Size, Yield, and Factor Intensity

As we seek to disentangle the different sources of dispersion in yield and input intensity, we next turn to a clue that emerges from a simple reduced form analysis of the data. Across all three countries, we observe a strong and consistent negative relationship between output per hectare and plot size. While this pattern is reminiscent of the long-standing discussion of an inverse farm size-yield correlation, we find in the final column of Tables 2 and 3 that this pattern holds across plots (planted with the same crop in the same season) within a farm. This differs from the usual finding that yield is negatively correlated with *farm size*. Across farms, factor market imperfections might explain an inverse relationship between land area and yield, but these market imperfections cannot explain this relationship across plots within a farm. Uganda exhibits the most extreme negative relationship between log yield and log plot size within a farm; the estimated elasticity is -0.75 (s.e.=0.02). Tanzania exhibits almost as strong a negative relationship, while in Ghana the estimated elasticity is -0.28 (s.e. = .04).

This pattern of a strong negative relationship between crop yields and plot size within a farm has been observed in multiple data sets from Africa (Carletto et al, 2015; Carletto et al., 2015; Bevis and Barrett, 2017). One source of this estimated inverse relationship might be measurement error in plot size. Kilic et al. (2016) provide a careful account of the role of this kind of measurement error using the same Uganda and Tanzania data sets that form part of our analysis. They show that while measurement error does contribute to the estimated inverse plot-size relationship, the relationship remains strong after using objective GPS measures of plot area and correcting for selection bias in the subset of plots measured with GPS.

Bevis and Barrett (2017) hypothesize that there is an “edge effect” on land productivity, in which plants near the boundary of a plot receive more attention in cultivation from the farmer, and perhaps have access to better nutrients and water then plants in the center of a plot. They provide evidence from the shape of plots that this edge effect explains part of the negative plot size – yield relationship in Uganda. Finally, Gourlay et al. (2017) report the results of a methodological experiment, also in Uganda, which carefully examined the problem of misreporting output data from farmers. They argue that farmers misreport crop harvests and that this measurement error is not random. Self-reported yields are biased upward compared to measurement of crop cuts at harvest. This effect is stronger on smaller plots. Gourlay et al. (2017) find that taking into account this measurement error fully explains the inverse plot size – yield relationship in a sample of farms in eastern Uganda.

We find, however, that labor per hectare is also strongly declining in plot size within a farm (Column 5 of Table 3). In Uganda and Tanzania, the elasticity of labor intensity with respect plot size is almost identical to that of yield with respect plot size. In Ghana, there is a much stronger negative relationship between labor intensity and plot size than there is between output and plot size, which may be accounted for by the positive relationship between capital per hectare and plot size reported in Table 4.

The consistency of the estimates of the correlations of plot size with labor intensity and with yield in Uganda and Tanzania may reflect similar patterns of measurement error in labor use as might exist in output. However, the Ghana results exhibit a different pattern, inconsistent with this parallel measurement error hypothesis.

An additional explanation for the inverse relationship between plot size and both labor intensity and yield is that smaller plots have higher unmeasured land quality.[16] If smaller plots are systematically better in terms of unobserved quality

[16] Barrett et al. (2010) argue against this hypothesis using data from Madagascar, showing that the introduction of a vector of objective measures of soil quality from soil tests has no effect on the estimated inverse yield – plot size relationship. However, the measures of land quality are not

than larger plots, households with a higher share of their land in smaller plots have higher unobserved wealth then otherwise similar households. We have detailed data on consumption in Uganda from a consumption module in the survey that is administered separately from the agricultural modules and therefore should be subject to different patterns of measurement error. In Column 1 of Table 5, we show (unsurprisingly) that consumption per adult equivalent is higher in Ugandan households that have more cultivated acreage. More importantly, we see that conditional on total cultivated area, households with more plots have higher consumption per adult equivalent. A household with one additional plot consumes on average 2% more per adult equivalent, conditional on total cultivated area.

In Column (2) of Table 5, we add a vector of household characteristics, including observed wealth and human capital, a vector of observed shocks, and measures of observed land quality. The relationship between the number of plots conditional on total cultivated area and consumption remains stable. In Column (3), we examine consumption changes over time by including a household fixed effect in a specification otherwise identical to that of Column (1), and the estimated coefficients on both total cultivated area and the number of plots fall and lose statistical significance. However, when we add back in Column (4) the vector of observed shocks used in Column (2), the estimated relationship between the number of plots and consumption regains its significance and most of its magnitude.

These correlations lead us to hypothesize that there is an important degree of unmeasured heterogeneity in land quality across the land of a given farmer. This is consistent with the patterns in Figures 1-3 documenting important dispersion in yield and factor intensity across the plots of an individual farmer. It may play a role in the strong inverse relationship observed between cultivation intensity and plot size across the plots of a farmer, and it is consistent with our finding of a positive correlation between household welfare and the number of plots cultivated, conditional on total area cultivated.

---

jointly significant predictors of yield, nor are they jointly significant in the production function estimated. This is a frequent characteristic of observed measures of land quality.

5. Theoretical framework

Our central argument is that heterogeneity in land quality plays an important role in explaining the dispersion of productivity at the level of plots and farms. This heterogeneity is unobservable to the econometrician but is well recognized by farmers. Some of the unobservables involve intrinsic properties of the soil or land, such as the physical and chemical properties of the soil, or the slope and topography. Other unobservables relate to highly localized shocks – such as hail that strikes one part of a farm but spares another part. Still others may involve complex interactions between location characteristics and plot characteristics: a heavy early-season rain makes one low-lying plot unworkable at the start of the season because of mud; but the same rainstorm is actually beneficial for another plot that is well drained. To help us understand the significance of this kind of location-specificity, we develop a model of agricultural production on heterogeneous land in which farmers can endogenously choose plot sizes and locations.

5.1 A model of agricultural production with endogenous plot selection

Let a household i hold a fixed endowment of land denoted by L_i. This land consists of a continuum of locations that can be indexed by k on the interval $[0, L_i]$.

At a location k, the quality of the land in effective units is denoted by $z_i(k)$. Anticipating an empirical implementation, we can think of $z(\cdot)$ as having an observable component and an unobservable component. Denote the observable portion of $z(\cdot)$ as $\eta(k)$ and let the unobservable portion be denoted by $u(k)$, such that $z(k) = \eta(k)\, u(k)$. For convenience, we assume that the functions $\eta(k)$ and $u(k)$ are both continuous across locations, so that the function $z_i(\cdot)$ is therefore continuous and integrable.

Land is used for producing an agricultural good. The production process uses a bundle of inputs that in principle could be applied on a location-specific basis. We denote the inputs used at a particular location as $X(k)$. Output is also affected by a location-specific productivity shock that depends on the state of the world, which we denote by $\gamma(k, s)$. The state of the world s is distributed according to

$\Delta(s)$ over support S. This shock is observed by the farmer before she chooses the input bundle. For example, this shock could consist of early-season rain – or perhaps the timing of the onset of the rainy season. A given state of the world may have different productivity implications for different locations on the farmer's land.

Given this notation, we define a simple production technology in which the output obtained by household i at location k conditional on the shock s having been realized will be given by:

$$Y_i(k,s) = \gamma_i(k,s) z_i(k) \left(X_i(k)\right)^{\theta}. \qquad (1)$$

If a profit-maximizing household were to farm only this single point, facing a household-specific shadow price w_i for inputs, the household would solve:

$$\max_{X_i(k)} \left[\gamma_i(k,s) z_i(k) \left(X_i(k,s)\right)^{\theta} - w_i X_i(k,s)\right] \qquad (2)$$

As an elementary optimality condition, this would give an optimum of $X_i^*(k,s) = \left(\frac{\theta\gamma_i(k,s)z_i(k)}{w_i}\right)^{\frac{1}{1-\theta}}$. The corresponding output would be:

$$Y_i^*(k,s) = z_i(k)\gamma_i(k,s)\left(\frac{\theta\gamma_i(k,s)z_i(k)}{w_i}\right)^{\frac{\theta}{1-\theta}}. \qquad (3)$$

Production could in principle be fine-tuned in this fashion to match the precise characteristics of each location, with inputs varying continuously across space. However, production takes place at the level of a *plot*. We define a plot (consistent with the definition used in most surveys) as a set of contiguous locations that are treated with an identical input bundle. In the rural contexts that we seek to model, the land is prepared in the same way and at the same time, the same crop (or crop mix) is planted across the plot, and the same inputs – including labor – are used across the plot.[17]

---

[17] We note that in some farm surveys, nomenclature varies. Our usage is consistent with the Uganda LSMS-ISA data, in which contiguous parcels of land are divided into plots. The individual plots may be planted with different crops and may be farmed with different inputs and management techniques.

For a household in our model, a plot will be defined as an interval $[\underline{k}, \overline{k}] \subseteq [0, L_i]$. The household faces a fixed cost c to create and farm a plot of land within its overall land holding. Because there is a fixed cost to farm a plot of land, there will be finitely many plots per farm. On a given plot j, the household i chooses an input vector x_{ij} that will be applied across all the locations in the plot. This input vector will depend on the realized shock s. The household chooses $X_{ij}(s)$ to solve:

$$\max_{X_{ij}(s)} \left[\left(X_{ij}(s)\right)^{\theta} \int_{\underline{k}}^{\overline{k}} \gamma_i(k,s) z_i(k)\, dk - w_i X_{ij}(s)\left(\overline{k} - \underline{k}\right) - c \right]. \qquad (4)$$

Let $z_{ij} = \int_{\underline{k}}^{\overline{k}} \gamma_i(k,s) z_i(k)\, dk$. This is essentially the weighted average, across locations, of the land productivity conditional on the realized productivity shock. Then the profit-maximizing input intensity at each point on the plot will be given by $X_{ij}^* = \left(\frac{\theta z_{ij}}{w_i}\right)^{\frac{1}{1-\theta}}$. Total input use will be simply $X_{ij}^*\left(\overline{k} - \underline{k}\right)$. Output will be $Y_{ij}^* = z_{ij}\left(\overline{k} - \underline{k}\right)\left(\frac{\theta z_{ij}}{w_i}\right)^{\frac{\theta}{1-\theta}}$. Note that the bluntness of input use means that the inputs applied at each location k will in general be different from the inputs applied if optimization were taking place separately at each location.

In general, the fixed cost of creating a plot implies that plots will have a minimum size. This implies that the profit-maximizing level of input use chosen for the whole plot will differ from that which would be chosen if the household were maximizing at each location separately. Output will differ correspondingly. The lone exception is the case in which the fixed cost $c \to 0$, in which case $\left(\overline{k} - \underline{k}\right) \to 0$ and $\left[Y_{ij}^* - \int_{\underline{k}}^{\overline{k}} Y_i^*(k) dk\right] \to 0$. With $c > 0$, the household chooses to divide its land into a finite number of plots. We next turn to the question of endogenous plot selection.

Consider first the household's option of producing on a single plot, $[0, L_i]$. The profit maximization problem is then given by:

$$\max_{X_i}\left[L_i X_i^{\theta}\int_0^{L_i}\gamma_i(k,s)z_i(k)\,dk-w_i x_i L_i-c\right]. \tag{5}$$

As an alternative to the single plot, the household could instead farm multiple plots. We assume that the household divides its landholding into plots at the start of the season, before inputs are chosen and – crucially – before the realization of the productivity shock $\gamma_i(k,s)$. In the model, this means that the boundaries of the plots are chosen before the shock is realized, and the inputs for each plot are chosen subsequent to this realization. In modelling the farm in this way, we seek to capture the notion that inputs can be adjusted throughout the growing season, so that the total input vector responds to the shocks. But plot boundaries cannot normally be adjusted once planting has taken place – and indeed, plot boundaries are often set even before planting, with a series of decisions that commit the household to planting certain crops at certain moments. For instance, the timing and techniques of land preparation will be linked to decisions about plot boundaries and potentially also crop choice.

Thus, we assume that the household chooses the boundaries between plots at a moment before the realization of the location-specific productivity shocks $\gamma_i(k,s)$.

Consider the problem of a household that is choosing a single boundary that will define two plots. Denote the threshold location between the two plots as L_{i1}, so that the two plots are $[0,L_{i1}]$ and $[L_{i1},L_i]$. In this case, an interior solution for the size of the two plots must hold; expected total profits could not be increased by moving this location either to the left or the right on the number line. The profit maximization problem can be written as:

$$\begin{aligned}\max_{L_{i1}}\int_{s\in S}\Bigg[\max_{X_{i1},X_{i2}}\Bigg[&L_{i1}X_{i1}^{\theta}\int_0^{L_{i1}}\gamma_i(k,s)z_i(k)\,dk+(L_i-L_{i1})X_{i2}^{\theta}\int_{L_{i1}}^{L_i}\gamma_i(k,s)z_i(k)\,dk\\&-w_iX_{i1}L_{i1}-w_iX_{i2}(L_i-L_{i1})-2c\Bigg]\Bigg]d\Delta(s).\end{aligned} \tag{6}$$

In effect, the household chooses the plot boundary L_{i1} to maximize expected profits, knowing what input bundle it would choose for each plot for every realization of the productivity shock $\gamma_i(k,s)$. The problem is well-defined.

Now consider a household that farms J plots, $J > 2$. We use the notation that L_{ij} will denote the right-hand boundary of the jth plot; i.e., the boundary between plot j and plot $j+1$. For notational convenience, we set $L_{i0} = 0$ and $L_{iJ} = L_i$. Then $\{L_{ij}\}_{j=0}^{J}$ is the sequence of plot boundaries. The first plot is given by the interval $[0, L_{i1}]$, and the jth plot covers the interval $[L_{ij-1}, L_{ij}]$, continuing to the Jth plot, which covers $[L_{iJ-1}, L_i]$.

We assume for convenience in what follows that all the plots are of sufficient quality that they will be actively farmed, allowing for an interior solution. The logic of the analysis would extend, however, to a situation in which the household chooses not to cultivate some portion of its land.

For notational convenience, let the size of the jth plot be denoted as $\tilde{L}_{ij} \equiv (L_{ij-1} - L_{ij})$. As before, the average productivity of plot j, conditional on the realization of the shock $\gamma_i(k,s)$, can be written as $z_{ij} = \int_{L_{ij-1}}^{L_{ij}} \gamma_i(k,s) z_i(k)\, dk$. Then the household's problem of choosing the boundaries of J plots can be written as:

$$E\hat{\pi}(J) = \max_{\{L_{ij}\}_{j=1}^{J}} \int_{s\in S} s \left[\max_{\{X_{ij}\}_{j=1}^{J}} \left[z_{ij}\tilde{L}_{ij}X_{ij}^{\theta} - \sum_{j=1}^{J} w_i X_{ij}\tilde{L}_{ij} - cJ \right] \right] d\Delta(s). \quad (7)$$

Because this problem is well-defined for any number of plots J, we can add an additional maximization over a finite number of possible values of J. Recall that for a single location k, the household can maximize profits conditional on the shock s with $Y_i^*(k,s) = z_i(k)\gamma_i(k,s)\left(\frac{\theta\gamma_i(k,s)z_i(k)}{w_i}\right)^{\frac{\theta}{1-\theta}}$. Let $\pi_i^*(k,s) = Y_i^*(k,s) - w_i X_i^*(k,s)$ and let $\pi_i^*(s) = \int_0^{L_i} \pi_i^*(k,s)\, dk$. This is the maximum profits that can be earned on the land with $c = 0$, conditional on the shock s. Then $\pi^* =$

$\int_{s\in S} \pi_i^*(k,s)\, d\Delta(s)$ is the expected maximum profits. Given this, $J^* = \left(\frac{\pi^*}{c} + 1\right)$ is an upper bound for the number of plots that can be profitably cultivated.

Thus, we can finally write $E\hat{\pi}^* = \max\ \{E\hat{\pi}(1),\ \ldots\ ,\ E\hat{\pi}(J^*)\}$.

In most of our empirical analysis, we will be concerned with input intensity on plots, rather than total input use. Note that our formulation lends itself well to this. Within any given plot, the optimization gives precisely an average input intensity and an average output per unit of land that are constant across the plot. As a result, at the plot level, we can carry out the analysis entirely in terms of input intensity and output per unit of land (i.e., yield).

We also note that as a simple extension of the analysis, we can let the input vector X be a Cobb-Douglas composite of two or more other inputs; e.g., labor N and chemicals V, such that $X = N^{\alpha}V^{1-\alpha}$. The analysis will go through unchanged.

Without imposing some further restrictions on the patterns of land quality, we cannot make any statements about the relationship between land quality and plot size. But we can offer a few relevant observations. First, we noted above that the maximum number of plots that could possibly be cultivated profitably by a household, $J^* = \left(\frac{\pi^*}{c} + 1\right)$. As pointed out above, this number depends inversely on the fixed cost. But J^* must also be positively related to the average land quality across the farm, which will enter into π^*. A farm household with very poor average land quality will *ceteris paribus* have a smaller maximum number of plots than a farm household with the same total land area but better quality land. This does not necessarily give rise to an empirical prediction, because farms will not in general cultivate the maximum possible number of plots. But it does point to an underlying pattern that should hold more generally: everything else equal, poor quality plots must be sufficiently large that they will earn positive profits.[18]

Consider the profit maximization from above for the jth plot cultivated by household i. We defined the size of this plot as $\tilde{L}_{ij} \equiv \left(L_{ij-1} - L_{ij}\right)$. The average

[18] Again, we assume for the moment that all land must be cultivated. We will address below the implications that arise when farms can leave land idle.

productivity of this plot, conditional on the realization of the shock $\gamma_i(k,s)$, is $z_{ij} = \int_{L_{ij-1}}^{L_{ij}} \gamma_i(k,s) z_i(k)\, dk$. We can solve the profit maximization problem and then ask, for a given value of z_{ij} what is the smallest plot size that will yield non-negative profits – or in other words, what threshold plot size will be needed to cover the fixed cost c. We can then ask how this plot size threshold changes in relation to z_{ij}. The formulation of this is straightforward. A plot with zero profits, in expectation, is one for which $z_{ij} X_{ij}^{\theta} \tilde{L}_{ij} - w_i X_{ij} \tilde{L}_{ij} = c$, or equivalently $\tilde{L}_{ij} = \frac{c}{z_{ij} X_{ij}^{\theta} - w_i X_{ij}}$. Substituting in the optimized value of X^*, we get a relationship between the threshold land quality that will sustain non-negative profits. The average plot quality cannot fall below z_{min}, which can be calculated as:

$$z_{min} = w^{\theta} c^{1-\theta} \left[\theta^{\frac{\theta}{1-\theta}} - \theta^{\frac{1}{1-\theta}}\right]^{\theta-1}. \qquad (8)$$

This relationship is illustrated in Figure 4, in which higher fixed costs for farming a plot are mapped into the minimum land quality needed to achieve zero profits. There will be no plots on which the average quality is lower than z_{min}. For plots of slightly higher quality, the plot size will need to be quite large to cover the fixed cost. Figure 5 shows how this minimum plot size will need to increase in relation to the average quality of the plot

This threshold plot size will also vary with the fixed cost. The higher the fixed cost of cultivating a plot, the larger will a plot needs to be, for any given quality, to make it profitably cultivable.

Within a farm, the optimal size of a plot depends on both the average quality of the land and the heterogeneity of the land quality. Holding average quality constant, the size of the plot will be decreasing in heterogeneity. Holding heterogeneity constant, the size of the plot will be decreasing in average quality (or put differently, it will increase on poorer land). The underlying logic is that there is a trade-off between the benefits gained by fine-tuning the inputs used on a plot, which tends to drive plot size smaller, and the fixed cost, which tends to drive plot size larger. On high-quality land, the fixed cost is a relatively smaller burden, and so plot size will be smaller, *ceteris paribus*. On low-quality land, the fixed cost poses a heavier burden, and so plot size will tend to be larger. At the

same time, however, the more heterogeneous a plot is in terms of land quality, the more costly it will be to have a large plot; a homogeneous plot can be large. In the extreme case of a farm that is entirely uniform in terms of land quality, there is no reason to subdivide this into plots, regardless of the quality.

6. Empirical analysis

This theoretical framework allows us to structure an empirical analysis of patterns of input intensities and yields across plots within a farm. The strong inverse relationship between plot size and cultivation intensity within a farm are both consistent with important permanent or transitory heterogeneity at small geographic scale in African agriculture.[19]

The importance of heterogeneity at small scale has implications for models of production in agriculture in Africa. In particular, there are implications for the role of factor market imperfections as sources of low measured productivity. Accounting for this local (and often unobservable) heterogeneity is also important for assessing the importance of misallocation.

Our goal in this section is to generate estimates of the relative importance of misallocation, risk, unobserved heterogeneity in land quality, and simple measurement error in generating the wide apparent dispersion of marginal productivities across farms and consequent losses in potential output.

The approach requires estimation of production functions, and we draw here on recent advances in the methods for identification of production functions, such as

[19] We are not the first to point to the importance of heterogeneity in agricultural systems at highly localized spatial scale. In the development economics literature, Hanna et al. (2014) show under experimental conditions the importance of localized variation. We also note that agricultural scientists have long recognized a high degree of localized heterogeneity in African agriculture. See, for example, Tittonell et al. (2005), Tittonell et al. (2007), Tittonell et al. (2008), Vanlauwe et al. (2006), Vanlauwe et al. (2015). The evidence suggests that this is not entirely a feature of African agriculture. Experimental work in other settings has shown very high degrees of spatial heterogeneity in soil quality even in large and seemingly uniform farm settings in the United States. See, for example, Hurley et al. (2004). The high degree of spatial heterogeneity in agriculture forms the basis for precision agriculture technologies, which are increasingly used in industrial agriculture; the basis and challenges for precision agriculture systems are surveyed in Stoorvogel et al. (2015). The economic returns to precision agriculture in the United States have recently been surveyed by Schimmelpfennig (2016).

the work done by Ackerberg et al. (2016), which in turn builds on Olley and Pakes (1992) and Levinsohn and Petrin (2003). This literature has suggested ways to address the chronic collinearity problems that have long bedeviled production function estimation. Although we do not claim to have completely solved the identification problem in the sections that follow, we are able to implement methods that address many of the most serious problems.

6.1 Analytic Framework

We begin by using the theoretical model of Section 4 to generate a framework for estimation. In keeping with our theory, suppose that farmer h produces output (Y_{hit}) on plot i in season t from land (L), labor (X) and capital (K) inputs according to

$$Y_{hit} = e^{\omega_{hit}}(L_{hit})^{\alpha_{Lhit}}(X_{hit})^{\alpha_{Xhit}}(K_{hit})^{\alpha_{Khit}}. \tag{9}$$

The parameter ω_{hit} is total factor productivity, which is at least partially unobserved to us, but which is known to the farmer. Given this structure, factor demands are subject to the classic production function endogeneity concern. In addition, unobserved heterogeneity in factor productivity gives us $\alpha_{Lhit}, \alpha_{Khit}$, and α_{Khit} that may be heterogeneous across plots.

We assume that farmers know the productivity of the factors they are using in cultivation, so factor demands will in general be correlated with the realizations of the factor productivities. Therefore, Equation (9) is an example of a model with correlated random coefficients.

Total factor productivity has three components. The first is a set of observable characteristics of the plot, farmer, or community, W_{Yhit}. The second is a component that is unobserved in the data but known to the farmer, ω_{Yhit}. Finally, there is a component that is unobserved in the data and unknown to the farmer at the time of input application, ϵ_{Yhit}. This final component could be actual output risk that is realized late in the season, or it could be pure measurement error in output; from the production function alone these cannot be distinguished.

$$\omega_{hit} = W_{Yhit}\beta_Y + \omega_{Yhit} + \epsilon_{Yhit} \tag{10}$$

Land, labor and capital inputs to production $(J \in \{L, X; K\})$ are modeled as the observed quantity of that factor (J^o_{hit}), observed as hectares, days, or value of

input J on plot i of household h in season t), corrected for a factor-specific set of observables (W_{Jhit}) and subject to classical measurement error ϵ_{Jhit}:

$$J_{hit} = J^{o}_{hit} e^{W_{Jhit}\beta_J - \epsilon_{Jhit}}. \tag{11}$$

The production function we estimate in logs, therefore is

$$\begin{aligned} y_{hit} = \alpha_{Lhit} l^{o}_{hit} + \alpha_{Xhit} x^{o}_{hit} + \alpha_{Khit} k^{o}_{hit} + W_{Yhit}\beta_Y \\ + \sum \alpha_{Jhit}\left(W_{Jhit}\beta_J - \epsilon_{Lhit}\right) + \omega_{Yhit} + \epsilon_{Yhit} \end{aligned} \tag{12}$$

The vector of observable determinants of total factor productivity ($W_{Yhit} = (W_{Ehit}, W_{Hhit})$) includes a rich set of indicators of shocks to productivity; most importantly, we have measures of the amount and timing of local rainfall interacted with characteristics of the plot and indicators of specific shocks (fire, flooding) on particular plots. We denote by W_{Ehit} the subset of those shock indicators that occur before the early harvest season begins, sufficiently early that farmers may be able to adjust factor inputs in response. Similarly, W_{Hhit} is the subset of those shock indicators that occur at harvest season, too late for farmers to adjust factor inputs in response.

Armed with this structure, we turn to the data from our three countries. Because the structures of agriculture (and the corresponding data sets) vary slightly across the countries, we consider the three countries separately.

6.2 Ghana

We begin with the Ghana sample, which were collected in the context of a randomized trial which provided farmers with substantial grants of crash and also varied the availability and cost of rainfall index insurance across farmers and over time.[20] These randomized treatments provide an important source of exogenous variation in factor use across farmers.

If factor markets are imperfect, then conditional on the realization of W_{Ehit} on plot i, the realizations of $W_{Eh,-i,t}$ on plots $-i \neq i$ of farmer h in season j provides variation in the shadow value of factors of production on plot i. Accordingly,

[20] See Karlan *et al*. (2014), section III and online appendix 1 for a detailed description of the sample, data collection procedures, index insurance and cash grants interventions and the randomization.

$W_{h,-i,t}$ and the treatment assignments comprise the set of instruments G_{hit} for plot i. Karlan et al. (2016) show that assignment to treatment strongly affects farmer demand for land and capital. Realizations of W_{Ehji} on plots of farmer h other than plot i strongly affect demand for all three factors of production. Table 6 presents IV estimates of $E(\alpha_J)$ under the assumption that the effects of Z_{hit} on $(l_{hit}, x_{hit}, k_{hit})$ are homogeneous (Heckman and Vytlacil 1998).[21] These coefficients imply a larger share of income for land than is observed in typical macroeconomic data, but recall that land in this exercise is defined as prepared plots and thus incorporates some share of either labor or capital in the form of tractor work.[22]

With estimates $\hat{\alpha}_J$ of the expected values of α_{Jhit} available, a first approximation to the distribution of log TFP across plots is simply:

$$\ln\widehat{TFP}^{a}_{hit} = y_{hit} - \hat{\alpha}_L l^{o}_{hit} - \hat{\alpha}_X x^{o}_{hit} + \hat{\alpha}_K k^{o}_{hit} - W_{Yhit}\hat{\beta}_Y - \sum \hat{\alpha}_J W_{Jhit}\hat{\beta}_J \quad (13)$$

Figure 7 provides an estimate of the density of $\ln\widehat{TFP}^{a}_{hit}$ in our Ghana sample. If we treat $\ln\widehat{TFP}^{a}_{hit}$ as if it were identical to the unobserved (to us) but known to the farmer total factor productivity that we defined as ω_{Yhit} in Equation (10), and if we furthermore neglect possible heterogeneity across plots in factor productivity and possible measurement error, then given $\hat{\alpha}_J$ and $\ln\widehat{TFP}^{a}_{hit}$, we can calculate the efficient allocation of factors across plots. The efficient allocation requires that for each factor J, $J^{e}_{hit} = s^{e}_{hit}\bar{J}$, where

$$\bar{J} = \frac{\sum_{\{h,i\}} J^{o}_{hit} e^{W_{Jhit}\hat{\beta}_J}}{N}$$

and

$$s^{e}_{hi} = \frac{\exp\left(\frac{1}{1-\hat{\alpha}_L-\hat{\alpha}_X-\hat{\alpha}_K}\ln\widehat{TFP}^{a}_{hit}\right)}{\frac{1}{N}\sum_{hi}\ln\widehat{TFP}^{a}_{hit}}$$

with N being the total number of plots. It is then possible to calculate the expected output of each farm given its estimated total factor productivity and the efficient allocation of factors to that plot. Table 7 summarizes this calculation. The first

[21] This is not a satisfactory assumption in this context, so the results should be treated as preliminary. Masten and Torgovitsky (2016) provide an estimator that is consistent for $E(\alpha_J)$ given the likely failure of the homogeneity assumption.

[22] In the United States a labor share is often taken to be about 50%, land perhaps 15% in capital about 35%.

row contains descriptive statistics on the observed distribution of output in the sample. The second row describes output conditional on an efficient reallocation of factors, given the estimated $\ln\widehat{TFP}_{hit}^{a}$. Total output increases by 546% upon this reallocation. The gain is large because the dispersion across plots in $\ln\widehat{TFP}_{hit}^{a}$ is very large and, as can be seen in Figures 8-10. there is virtually no correlation between $\ln\widehat{TFP}_{hit}^{a}$ and the observed allocation of factors across plots in the sample. This is very similar to the finding of Restuccia and Santaeulalia-Llopis (2017), although their method of estimating TFP dispersion is quite different from ours.[23]

However, there are three causes for concern regarding this first estimate of the extent of factor misallocation. First, $\ln\widehat{TFP}_{hit}^{a}$ is an overestimate of ω_{Yhit}. From (12) and (13), we observe:

$$\ln\widehat{TFP}_{hit}^{a} = \underbrace{\omega_{Yhit}}_{\text{unobs tfp}} + \underbrace{\sum_{J}(\alpha_{Jhit}-\hat{\alpha}_J)(W_{Jhit}\hat{\beta}_J)}_{\substack{\text{unobs productivity of}\\ \text{observed characteristics}}} + \underbrace{\sum_{J}(\alpha_{Jhit}-\hat{\alpha}_J)J_{hit}}_{\substack{\text{unobs productivity}\\ \text{of factors}}}$$
$$- \underbrace{\sum_{J}(\alpha_{Jhit}-\hat{\alpha}_J)\epsilon_{Lhit}}_{\substack{\text{unobs productivity of}\\ \text{factor measurement error}}} - \underbrace{\sum_{J}\hat{\alpha}_J\epsilon_{Jhit}}_{\substack{\text{factor}\\ \text{measurement error}}} \tag{14}$$
$$+ \underbrace{\epsilon_{Yhit}}_{\substack{\text{post-input}\\ \text{shocks measurement}\\ \text{error in y}}}$$

The final three terms are sources of variation in *measured* productivity (i.e., $\ln\widehat{TFP}_{hit}^{a}$), but they do not give rise the kinds of *actual* productivity variation that could be remedied through reallocation. Reallocation cannot "solve" measurement error, nor can reallocation "solve" late-season idiosyncratic shocks that affect yield.

A second concern regarding our estimate of the impact of misallocation is that heterogeneity complicates the notional undistorted benchmark. To be specific, the

[23] In particular, they assume that the parameters of the agricultural production function are known, and they use this to back out farm-level TFP from the observed inputs and outputs.

unobserved variation in factor productivity α_{Jhit} implies that the efficient allocation of factors does not equalize factor proportions across plots. Thus, our reallocation exercise is based on a slightly flawed benchmark.

Third, classical additive measurement error in log factor inputs induces an overestimate of average factor use because of the concavity of the log function.

Our estimation approach also allows us to estimate the density of $\ln\widehat{TFP}^{a}_{hit}$ across plots within farms. This too is shown in Figure 7. The dispersion of within-farm productivity is surprisingly large, relative to that across the Northern Region of Ghana as a whole. One might expect that farmer effects would be an important component of Total Factor Productivity, and similarly that factor market imperfections might generate distortions in factor ratios across farms that would be misinterpreted as TFP differences. However, the wide dispersion within farms raises the possibility that there is extremely high local variation in unobserved plot-level productivity. Alternatively, measurement error and late season shocks could be particularly important. What is clear is that the high dispersion in productivity across farms is mirrored by relatively high dispersion in productivity across plots within farms. This suggests that differences in farmer quality may not account well for the data.

Without knowledge of the frictions generating factor market misallocation, it is difficult to disentangle the variation in unobserved total factor productivity ω_{Yhit} from the variation in factor-specific productivities α_{Jhit}, measurement error, or late season shocks. However, an efficient allocation of factors across plots cultivated by the same farmer implies patterns of behavior that can identify the variances and covariances of these productivity shocks and measurement errors. We documented above that there is significant variation in yields and factor intensity across plots within a farm. If factors are allocated efficiently across these plots, then the observed variation is being driven by some combination of these unobserved productivity shocks and measurement error. For example, if we observe variation in labor inputs that is not correlated with output or other inputs, then our model would attribute this to measurement error in labor.

In what follows, we maintain the assumption that the allocation of factors within a farm is efficient. By the second welfare theorem, there exist farm-season-specific

shadow prices $(p_{Lht}, p_{Xht}, p_{Kht})$, with $p_{Yht} = 1$ the numeraire, such that the input choices maximize plot profits. Therefore, actual (not observed) factor inputs satisfy

$$p_{Lht} = \frac{\alpha_{Lhit} Y_{hit} e^{-W_{Hhit}\beta_H - \epsilon_{Yhit}}}{L_{hit}}; p_{Xht} = \frac{\alpha_{Xhit} Y_{hit} e^{-W_{Hhit}\beta_H - \epsilon_{Yhit}}}{X_{hit}}; \\ p_{Kht} = \frac{\alpha_{Khit} Y_{hit} e^{-W_{Hhit}\beta_H - \epsilon_{Yhit}}}{K_{hit}}, \tag{15}$$

in familiar Cobb-Douglas fashion. We specify the factor-specific productivity shocks as

$$\alpha_{Jhit} = e^{\alpha_J + \omega_{Jhit}}$$

and because the model incorporates Total Factor Productivity shocks as well, we normalize the factor-specific shocks along the simplex:

$$\omega_{Lhit} + \omega_{Xhit} + \omega_{Khit} = 0.$$

Solving for the reduced forms, taking logs and substituting observed variables for unobserved true factor inputs, we find that log output is:

$$\begin{aligned} y_{hit} = W_{Hhit}\beta_H + \epsilon_{yhit} \\ &+ \frac{1}{1 - \sum_J \alpha_{Jhit}} \left\{ \frac{W_{Ehit}\beta_E}{1 - \sum_J \alpha_{Jhit}} + \omega_{Yhit} \right. \\ &+ \alpha_{Lhit} \ln\left(\frac{\alpha_{Lhit}}{p_{Lht}}\right) + \alpha_{Xhit} \ln\left(\frac{\alpha_{Xhit}}{p_{Xht}}\right) \\ &\left. + \alpha_{Khit} \ln\left(\frac{\alpha_{Khit}}{p_{Kht}}\right) \right\} \\ &= W_{Hhit}\beta_H + \epsilon_{yhit} + q_{hit} \end{aligned} \tag{16}$$

In this formulation, q_{hit} is the TFP of plot hit, inclusive of the factor-specific productivities. Factor demand is determined by

$$\begin{aligned} l^o_{hit} &= \alpha_L + \omega_{Lhit} - \ln(p_{lht}) - W_{Lhit}\beta_L + \epsilon_{Lhit} + q_{hit} \\ x^o_{hit} &= \alpha_X + \omega_{Xhit} - \ln(p_{Xht}) - W_{Xhit}\beta_X + \epsilon_{Xhit} + q_{hit} \\ k^o_{hit} &= \alpha_K + \omega_{Khit} - \ln(p_{Kht}) - W_{Khit}\beta_K + \epsilon_{Khit} + q_{hit} \end{aligned} \tag{17}$$

Observed factor inputs, unsurprisingly, depend on the cost of that factor, on the realization of that factor's productivity and total productivity, and the realization of that factor's measurement error.

Factor shadow costs vary in unspecified ways across farmers as a consequence of factor and other market imperfections. Therefore, we focus on within-farm cross-plot variation, taking farmer-year fixed effects $(\lambda_{Yht}, \lambda_{Lht}, \lambda_{Xht}, \lambda_{Kht})$ out of (16) and (17). Given estimates of $\hat{\alpha}_J$ and $\hat{\beta}_J$, we calculate

$$\begin{aligned}
\tilde{y}_{hit} &\equiv y_{hit} - W_{Hhit}\hat{\beta}_H = \lambda_{Yht} + \epsilon_{yhit} + q_{hit} \\
\tilde{l}_{hit} &\equiv l^o_{hit} - W_{Lhit}\hat{\beta}_L = \lambda_{Lht} + \omega_{Lhit} + \epsilon_{Lhit} + q_{hit} \\
\tilde{x}_{hit} &\equiv x^o_{hit} - W_{Xhit}\hat{\beta}_X = \lambda_{Xht} + \omega_{Xhit} + \epsilon_{Xhit} + q_{hit} \\
\tilde{k}_{hit} &\equiv k^o_{hit} - W_{Khit}\hat{\beta}_K = \lambda_{Kht} + \omega_{Khit} + \epsilon_{Khit} + q_{hit}
\end{aligned} \tag{18}$$

The observable within-farm variation in output, factor inputs, factor ratios and output per factor input then provides us with sufficient information to disentangle the roles of measurement error and post-input risk from those of productivity shocks in the observed cross-plot, within-farm variation we described in Section 3.

Table 8 lists the 14 variances and covariances that we can identify from the observed within-farm variances.

Variances of factor ratios and output per factor across plots within farms provide information on factor-specific shocks and measurement error in factors and output.

$$\begin{aligned}
var(\tilde{x}_{hit} - \tilde{l}_{hit}) &= \sigma_L^2 + \sigma_X^2 - 2\sigma_{LX} + \sigma_{\epsilon L}^2 + \sigma_{\epsilon X}^2 \\
var(\tilde{k}_{hit} - \tilde{l}_{hit}) &= \sigma_L^2 + \sigma_K^2 - 2\sigma_{LK} + \sigma_{\epsilon L}^2 + \sigma_{\epsilon K}^2 \\
var(\tilde{x}_{hit} - \tilde{k}_{hit}) &= \sigma_X^2 + \sigma_K^2 - 2\sigma_{XK} + \sigma_{\epsilon X}^2 + \sigma_{\epsilon K}^2 \\
var(\tilde{y}_{hit} - \tilde{l}_{hit}) &= \sigma_L^2 + \sigma_{\epsilon y}^2 + \sigma_{\epsilon L}^2 \\
var(\tilde{y}_{hit} - \tilde{x}_{hit}) &= \sigma_X^2 + \sigma_{\epsilon y}^2 + \sigma_{\epsilon X}^2 \\
var(\tilde{y}_{hit} - \tilde{k}_{hit}) &= \sigma_K^2 + \sigma_{\epsilon y}^2 + \sigma_{\epsilon K}^2
\end{aligned} \tag{19}$$

Variances of output and inputs provide information on variation in TFP across plots:

$$\begin{aligned}
var(\tilde{y}_{hit}) &= \sigma_Q^2 + \sigma_{\epsilon Yhit}^2 \\
var(\tilde{l}_{hit}) &= \sigma_L^2 + \sigma_{\epsilon L}^2 + \sigma_Q^2 + 2\sigma_{QL} \\
var(\tilde{x}_{hit}) &= \sigma_X^2 + \sigma_{\epsilon X}^2 + \sigma_Q^2 + 2\sigma_{QX} \\
var(\tilde{k}_{hit}) &= \sigma_K^2 + \sigma_{\epsilon K}^2 + \sigma_Q^2 + 2\sigma_{QK}
\end{aligned} \tag{20}$$

Finally, the normalization of factor specific productivities distinguishes these from TFP shocks and measurement error:

$$\begin{aligned}
\sigma_K^2 &= \sigma_L^2 + \sigma_X^2 + 2\sigma_{LX} \\
\sigma_{LK} &= -\sigma_L^2 - \sigma_{LX}
\end{aligned} \tag{21}$$

$$\sigma_{XK} = -\sigma_X^2 - \sigma_{LX}$$
$$\sigma_{QK} = -\sigma_{QL} - \sigma_{QX}$$

These restrictions exactly identify the underlying variances across plots within farms. In particular, they provide estimates of the variances across plots within farms of measurement error in output or post-input shocks ϵ_{Yhit}, measurement error in each of the observed factors ϵ_{Jhit}, and of factor-specific productivity shocks ω_{Jhit}. These estimates are provided in Table 8.

There is little evidence of variation in unobserved TFP across plots within a farm. It does appear that factor-specific productivities vary widely across plots of farms, and classical measurement error is a still more important source of variation. The measurement error is largest for capital, and smallest for land. The estimates indicate a strong negative covariance between productivity shocks to labor and capital which may have to do with the intuitive notion that these factors are strong substitutes in African agriculture. Finally, the estimates indicate that there is a high level of either measurement error in output or a late-season production risk across plots within farms.

We have been looking at evidence on the within-farm across-plot distributions of productivity shocks and measurement errors. Next we turn to distribution across farms. This requires some homogeneity assumption. It would clearly be an error to assume that plot-specific TFP (q_{hit}) would be distributed across farms in the same way that it is distributed within a farm. One of the components of q_{hit} is ω_{hit}, which likely contains a persistent farmer-level component that is absorbed by the farm fixed effect in our procedure. The factor-specific productivity components ω_{Jhit} may also contain a farmer-specific component. Instead, we proceed by beginning with the first approximation to measuring total factor productivity, $\ln \widehat{TFP}_{hit}^a$, and shrinking its variance to account for the variances of the measurement errors and late-season shocks we can measure within the farm.

The final two terms of equation (14) represent components of $\ln \widehat{TFP}_{hit}^a$ that are not susceptible to efficient reallocation, and they are uncorrelated with the remaining terms. We have estimates of the within-farm variances of both ϵ_{Yhit} (late season shocks and output measurement error) and $\sum \alpha_J \epsilon_{Lhit}$ (factor

measurement error).[24] We construct a revised estimate of plot-level total factor productivity with the assumption that the overall variances of $\sum \alpha_J \epsilon_{Lhit}$ and ϵ_{Yhit} are the same as these within-farm variances:

$$\ln \widehat{TFP}^b_{hit} = \ln \widehat{TFP}^a_{hit} * \left(\frac{var\left(\ln \widehat{TFP}^a_{hit} \right) - \sigma^2_{\epsilon Y} - \sum_J (e^{\alpha_j})^2 \sigma^2_{\epsilon J}}{var\left(\ln \widehat{TFP}^a_{hit} \right)} \right)^{\frac{1}{2}} \qquad (22)$$

If there are farmer-specific components to either late season shocks or to measurement error, then our within farm estimates of these variances are conservative and the associated adjustment to $\ln \widehat{TFP}^a_{hit}$ is moderated.

Figure 7 provides a kernel estimate of the density of $\ln \widehat{TFP}^b_{hit}$. It is immediately apparent that accounting for measurement error in factors and output and for late-season shocks dramatically reduces the apparent dispersion of TFP across plots.

Finally, we turn to the gains from optimal reallocation of factors across plots. Our hypothetical reallocation rule assumes that an efficient social planner could (costlessly) reallocate capital, labor and land across the plots of different farmers subject to not exceeding the total stock of these factors of production. We attribute total factor productivity to the farmer-plot. The reallocation accomplishes two objectives: first, it equalizes the ratios of marginal products of the different factors across all plots: second it allocates more of all factors to the plots with higher productivity. Farmers with high TFP plots expand those plots at the expense of farmers with low TFP plots. However, a portion of the observed dispersion in commonly-measured “TFP” across plots is due to measurement error and to risk; since this is not actually productivity dispersion, it is not susceptible to reallocation. In short, we compare the “true” productivity across plots in northern Ghana, after factors are reallocated, with the “true” productivity before reallocation (as opposed to the measured productivity, which is distorted by measurement error and by late season risk.)

---

[24] We cannot separately identify the contribution to the variance of $\widehat{tfp}^a_{hit}$ attributable to $\sum(\alpha_{Jhit} - \alpha_J)\epsilon_{Lhit}$ – the interaction between factor measurement error and unobserved variation in factor productivity, because the covariance of this term with the remainder of $\widehat{tfp}^a_{hit}$ depends on the covariances of α_{Jhit} and W_{Jhit} and observed factor inputs, which we do not know. However, as long as existing factor allocations are not negatively correlated with unobserved factor productivity, this term makes a positive contribution to the overall variance.

The third row of Table 7 reports the results of optimally reallocating factors across plots in the Northern Region after taking into account the existence of late-season risk and measurement error. The gains remain substantial; approximately 40%. These are much smaller, however, than the estimates that we would have calculated using the residuals of the production function.

There is one final adjustment to be made to this first approximation of the gain to reallocation. Actual factor inputs are:

$$J_{hit} = J_{hit}^{o} e^{W_{Jhit}\beta_J - \epsilon_{Jhit}}$$

If ϵ_J is approximately normal, the expected value of factor inputs conditional on observed inputs is:

$$E(J_{hit}|J_{hit}^{o}) = J_{hit}^{o} e^{W_{Jhit}\beta_J} E(e^{-\epsilon_{Jhit}}) = J_{hit}^{o} e^{W_{Jhit}\beta_J} e^{-\frac{1}{2}\sigma_J^2}.$$

The stock of factors available for reallocation is somewhat smaller than would be the case if there were not measurement error, because with mean 1 multiplicative measurement error in levels (classical measurement error in logs), the absolute magnitude of the positive errors is larger than that of the negative errors. This correction is entirely a consequence of the assumed form of the measurement errors in factor inputs; therefore we separate out its effect on the estimated magnitude of the gains to reallocation. Taking this correction into account reduces the potential gain from redistribution still further, to 15%, as shown in the bottom row of Table 7.

Our analysis for Ghana thus shows that the gains from reallocation are overestimated in an analysis based simply on the residuals from the production function. The same concern would apply to an approach based on a calibrated model. In both cases, the residual (or farm-level TFP estimate) is effectively forced to capture a number of sources of productivity dispersion that are not susceptible to reallocation. For instance, these approaches treat shocks and measurement error as though they are part of the true productivity of farms and farmers. When we make appropriate adjustments, we find that these shocks and mismeasurement are important in explaining the observed dispersion of productivity. We also find that heterogeneity in plot-specific TFP is large, and this

too is not susceptible to reallocation. For Ghana, taking into account this fuller picture of productivity dispersion leads us to much smaller estimates of the costs of misallocation (or equivalently the gains from reallocation). We find that optimal reallocation across the entire Northern Region would give an increase in productivity of 15%, more than an order of magnitude smaller than the estimates obtained without correct adjustment.

6.3 Uganda

We turn next to data from Uganda. The Uganda data set has the advantage of being a large panel that tracks parcels (the contiguous area from which individual plots are carved) over six seasons (three years, with two growing seasons in each year). As in Ghana, we have rich information on local weather and rainfall variation that affects plot productivity differently depending upon their characteristics. We make use of shocks on the farms of other farmers in the community as across an instrument for inputs of a farmer, and of shocks on the other plots of a given farmer as an instrument for plot level inputs of the farmer. However, there is no explicit randomization that we can use to identify variation in shadow costs of factors across farmers. We therefore adopt a different empirical strategy for production function estimation by making use of the dynamics of the productivity process.

The production function in Uganda is simpler than that in Ghana, because there are no significant capital inputs or purchased intermediate goods.[25] Hence, output on plot i of farmer h in season t is

$$Y_{hit} = e^{W_{hit}\beta+\epsilon_{hit}}e^{\omega_{hit}}(L_{hit})^{\alpha_L}\ (X_{hit})^{\alpha_X} \qquad (23)$$

where L and X are the actual land and labor inputs. W_{hit} is a vector of observed determinants of productivity, such as land characteristics, weather outcomes and plot-specific shock variables. The timing of the realization of these shocks will be important below, and it varies across the different elements of W_{hit}.

Unobserved total factor productivity is

$$\omega_{hit} = \lambda_h + \lambda_{hi} + \zeta_{ht} + \zeta_{hit}$$

with ζ_{ht} and ζ_{hit} being productivity shocks that are assumed to be iid over time. They are realized after the allocation of land is fixed, but with labor still flexible. Farmer ability (λ_h) and plot-specific productivities (λ_{hi}) are known to the farmer.

[25] For example, fewer than 3 percent of households use chemical fertilizers.

We cannot distinguish between measurement error in output and late-season shocks to output; both are captured by ϵ_{hit}. Land and labor may be measured with error as well, with $L_{hit} = L^{o}_{hit} e^{-\epsilon_{Lhit}}$ and $X_{hit} = X^{o}_{hit} e^{-\epsilon_{Xhit}}$. Taking all of these together and using logs, (23) becomes

$$y_{hit} = \alpha_X x_{hit} + \alpha_L l_{hit} + W_{hit}\beta + \lambda_h + \lambda_{hi} + \zeta_{ht} + \zeta_{hit} - \alpha_X \epsilon_{xhit} - \alpha_L \epsilon_{lhit} + \epsilon_{iht} \tag{24}$$

We define Z_{ijt} as a vector of shocks that affect labor input this period, or labor or land input last season, but that do not directly affect output or measurement error. The key elements of $\tilde{Z}_{ijt}$ are subsets of the shocks in $W_{h,-i,t}$ and $W_{-h,it}$ at time t and t-1; that is, shocks on other plots of farmer h and on the plots of other farmers.

In log first differences, (24) becomes

$$\begin{aligned} y_{hit} - y_{hit-1} = {} & \alpha_X (x_{hit} - x_{hit-1}) + \alpha_L (l_{hit} - l_{hit-1}) \\ & + (W_{hit} - W_{hit-1})\beta + \zeta_{ht} + \zeta_{hit} - \zeta_{ht-1} - \zeta_{hit-1} \\ & - \alpha_X (\epsilon_{xhit} - \epsilon_{xhit-1}) - \alpha_L (\epsilon_{lhit} - \epsilon_{lhit-1}) + \epsilon_{iht} \\ & - \epsilon_{iht-1} \end{aligned} \tag{25}$$

We estimate (25) via GMM using 2 key moment conditions. The first is from (23) and relies on weather and other plot shocks Z_{hit}; the second is from (24) and adds twice lagged inputs as valid instruments:

$$\begin{aligned} & E[(\lambda_h + \lambda_{hi} + \zeta_{ht} + \zeta_{hit} - \alpha_X \epsilon_{xhit} - \alpha_L \epsilon_{lhit} + \epsilon_{iht}) \cdot (Z_{hit})] \\ & E\left[\begin{pmatrix} \zeta_{ht} + \zeta_{hit} - \zeta_{ht-1} - \zeta_{hit-1} + \theta(\epsilon_{xhit} - \epsilon_{xhit-1}) \\ +\gamma(\epsilon_{lhit} - \epsilon_{lhit-1}) + \epsilon_{iht} - \epsilon_{iht-1} \end{pmatrix} \cdot \begin{pmatrix} l_{hit-2} \\ x_{hit-2} \\ Z_{hit} \end{pmatrix}\right] \end{aligned} = 0 \tag{26}$$

GMM estimates of the production function are presented in Table 9. As in Ghana, we can use the empirical residuals from equation (23) to construct a first approximation to total factor productivity. Figure 11 provides an estimate of the density of $\widehat{\ln TFP}^{a}_{hit}$ in our Uganda sample. The estimated density is dramatically more dispersed in Uganda then it was in Ghana; perhaps again reflecting the greater diversity of farming conditions across Uganda as a whole than across the Northern Region of Ghana.

If we treat $\widehat{\ln TFP}^{a}_{hit}$ as if were identical to the "true" TFP (unobserved by us but known to the farmer) and if we neglect possible measurement error, then given $\hat{\alpha}_J$

and $\widehat{\ln TFP}^{a}_{hit}$, we can calculate the efficient allocation of factors across plots, using the same procedure as for Ghana. The first row of Table 10 provides summary statistics on the actual harvests in Uganda; the second row provides similar statistics on the hypothetical harvest expected with an optimal reallocation of factors from low-TFP to high-TFP and to with efficient levels of labor relative to land. As in Ghana, the gain to an efficient reallocation of factors is extremely high. Output could be expected to be 800% higher than we observe in the data. This is a consequence both of the wide dispersion of $\widehat{\ln TFP}^{a}_{hit}$ and the fact that the allocation of labor and land across plots is virtually uncorrelated with $\widehat{\ln TFP}^{a}_{hit}$, as shown in Figures 12 and 13.

However, $\widehat{\ln TFP}^{a}_{hit}$ incorporates measurement error in land, labor, and output as well as late-season shocks to output, none of which provide sources of productivity variation that could permit gains through factor reallocation. In addition, we explore here the possibility that adjustment of labor input is more flexible than that of land inputs, based on the sequence of operations through an agricultural cycle.

We again make use of the claim that the allocation of factors across plots *within* a farm is efficient, and thus rely on the existence of farm-season-specific shadow prices at which the observed allocation of factors maximizes profits. We divide $W_{hit} \equiv (W_{Ehit}, W_{Hhit})$ into early- and harvest-season observable plot characteristics and recall that ζ_{ht}, ζ_{hit} are revealed after land area is chosen. These assumptions generate the following reduced forms for input demand and output:

$$\begin{aligned} x_{hit} &= F_{xht} + \frac{W_{Ehit}\beta_E}{(1-\alpha_L-\alpha_X)} + \frac{W_{Hhit}\beta_H + \zeta_{hit}}{(1-\alpha_X)} + \frac{\lambda_{hi}}{(1-\alpha_L-\alpha_X)} \\ &\quad + \epsilon_{xhit} \\ l_{hit} &= F_{lht} + \frac{W_{Ehit}\beta_E + \lambda_{hi}}{1-\alpha_L-\alpha_X} + \epsilon_{hlit} \\ y_{hit} &= F_{yht} + \frac{W_{Ehit}\beta_E}{(1-\alpha_L-\alpha_X)} + \frac{W_{Hhit}\beta_H}{(1-\alpha_X)} + \frac{\zeta_{hit}}{(1-\alpha_X)} \\ &\quad + \frac{\lambda_{hi}}{(1-\alpha_L-\alpha_X)} + \epsilon_{hit} \end{aligned} \tag{27}$$

The farmer-year fixed effects reflect variation across farmer-years in the shadow values of land, labor and output, along with any permanent variation in productivity across farmers. We look at variation within farmer-years, differencing farmer-year means from plot level data: $dx_{it} = x_{it} - x_{ht}$, and analogously for all variables.

Define:

$\widetilde{dx_{hit}} = dx_{hit} - \frac{1}{(1-\alpha_L-\alpha_X)}(dW_{Ehit})\,\beta_E - \frac{1}{(1-\alpha_X)}(dW_{Hhit})\beta_H,$

$\widetilde{dl_{hit}} = dl_{hit} - \frac{1}{1-\alpha_L-\alpha_X}(dW_{Ehit})\beta_E$, and

$\widetilde{dy_{hit}} = dy_{hit} - \frac{1}{(1-\alpha_L-\alpha_X)}(dW_{Ehit})\beta_E - \frac{1}{(1-\alpha_X)}(dW_{Hhit})\beta_H.$

Efficient allocation of factors across plots within farms implies some specific relationships between variances and covariances observed in the data and the underlying distributions of measurement error and productivities across plots within farms. As in Ghana, we cannot identify the variance of the farmer-level component of productivity shocks using within-farm variation. So instead, our goal is to estimate the level of measurement error in factor inputs $(\sigma^2_{\epsilon l}, \sigma^2_{\epsilon x})$ and combined measurement error and late season productivity shock to output $(\sigma^2_{\epsilon y})$. We will also be able to identify the variances of cross-plot permanent and transitory productivity $\left(\sigma^2_{\lambda i}, \sigma^2_{\zeta i}\right)$. We use:

$$cov(\widetilde{dy}, \widetilde{dx}) = \left(\frac{1}{(1-\alpha_L-\alpha_X)}\right)^2 \sigma^2_{\lambda i} + \left(\frac{1}{(1-\alpha_X)}\right)^2 \sigma^2_{\zeta i}$$

$$cov(\widetilde{dl}, \widetilde{dx}) = \frac{1}{(1-\alpha_L-\alpha_X)^2}\sigma^2_{\lambda i}$$

$$var(\widetilde{dl}) = \sigma^2_{\epsilon l} + \left(\frac{1}{1-\alpha_L-\alpha_X}\right)^2 \sigma^2_{\lambda i} \tag{28}$$

$$var(\widetilde{dx}) = \left(\frac{1}{(1-\alpha_L-\alpha_X)}\right)^2 \sigma^2_{\lambda i} + \left(\frac{1}{(1-\alpha_X)}\right)^2 \sigma^2_{\zeta i} + \sigma^2_{\epsilon x}$$

$$var(\widetilde{dy}) = \left(\frac{1}{(1-\alpha_L-\alpha_X)}\right)^2 \sigma^2_{\lambda i} + \left(\frac{1}{(1-\alpha_X)}\right)^2 \sigma^2_{\zeta i} + \sigma^2_{\epsilon y}$$

Estimates of these variances are provided in Table 11. Most notable are the small (relative to Ghana) magnitudes of measurement error in land and labor.

Interestingly, the magnitude of measurement error/late season shocks seem similar in Ghana and Uganda.

As in Ghana, we construct a second estimate of log TFP by shrinking $\ln\widehat{TFP}_{hit}^{a}$ to account for measurement error:

$$\ln\widehat{TFP}_{hit}^{b} = \ln\widehat{TFP}_{hit}^{a} * \left(\frac{var\left(\ln\widehat{TFP}_{hit}^{a}\right) - \sigma_{\epsilon Y}^{2} - \alpha_{L}^{2}\sigma_{\epsilon L}^{2} - \alpha_{X}^{2}\sigma_{\epsilon X}^{2}}{var\left(\ln\widehat{TFP}_{hit}^{a}\right)}\right)^{\frac{1}{2}} \quad (29)$$

An estimate of the density of this second estimate of log TFP is provided in Figure 11.

We will use these estimates to recalculate the potential gains to reallocation. We correct the expected value of factor inputs for measurement error as we did for Ghana. In addition, we take account of the presumption that land allocation is decided before the realization of ζ_{hit}. To do so, we generate

$$\hat{\zeta}_{hit} = \left(\frac{\hat{\sigma}_{\zeta}^{2}}{var\left(\ln\widehat{TFP}_{hit}^{a}\right)}\right)^{\frac{1}{2}} \ln\left(\widehat{TFP}_{hit}^{a}\right) \quad (30)$$

and calculate the optimal reallocation subject to the constraint that land allocations cannot respond to this shock.[26]

Summary statistics on the gains from this optimal reallocation are provided in the third row of Table 10. The expected gain in output over the existing allocation is 61 percent. As in Ghana, the expected gains to reallocation are far smaller after taking into account the effects of measurement error and late season shocks on estimates of total factor productivity.

[26] We assume that the variance of ζ_{ht}, which we cannot estimate, is zero. This maximizes the potential gain from reallocation. The optimal allocation of land is

$$\frac{L_{hit}}{L} = \frac{e^{\frac{\left(W_{Ehit}\beta_{E} + \ln\widehat{TFP}_{hit}^{b} - \hat{\zeta}_{hit}\right)}{1-\alpha_{L}-\alpha_{X}}}}{\sum e^{\frac{\left(W_{Ehit}\beta_{E} + \ln\widehat{TFP}_{hit}^{b} - \hat{\zeta}_{hit}\right)}{1-\alpha_{L}-\alpha_{X}}}}$$

And for labor:

$$\frac{e^{\frac{\left(W_{Ehit}\beta_{E} + \ln\widehat{TFP}_{hit}^{b} - \hat{\zeta}_{hit}\right)}{1-\alpha_{L}-\alpha_{X}} + \frac{\left(W_{Hhit}\beta_{H} + \hat{\zeta}_{hit}\right)}{1-\alpha_{X}}}}{\sum e^{\frac{\left(W_{Ehit}\beta_{E} + \ln\widehat{TFP}_{hit}^{b} - \hat{\zeta}_{hit}\right)}{1-\alpha_{L}-\alpha_{X}} + \frac{\left(W_{Hhit}\beta_{H} + \hat{\zeta}_{hit}\right)}{1-\alpha_{X}}}}$$

Given the large geographical dispersion across Uganda, rows 3 and 4 provide calculations for the gains to efficient reallocations across farms within a region and within a village, using $\ln \widehat{TFP}_{hit}^{a}$ and $\ln \widehat{TFP}_{hit}^{b}$.

7. Discussion

The results from Ghana and Uganda show the importance of accounting carefully for measurement error, shocks, and heterogeneity in technology (including input quality) in measuring productivity at the level of individual production units. These issues have previously been raised in critiques of the literature on misallocation, but the data from African farms provide sufficiently rich detail that we can begin to disentangle the different sources of productivity dispersion. Our analysis suggests that previous estimates of misallocation have probably overestimated the potential productivity losses due to misallocation (or, equivalently, the gains from efficient reallocation). We do find that the gains from reallocation are non-trivial, but they are certainly not of such a magnitude as to account in a macro sense for the aggregate differences in agricultural productivity – or income per capita – between rich and poor countries.

Given that reallocation would also entail massive costs – not least, in terms of the social welfare implications of reallocating land away from many poor smallholders in Africa – we believe that these findings are important for their own sake. But in addition, we believe that there are additional implications for the broader literature that has grown up around the topic of misallocation in development and growth. Much of this literature has relied on cross-section data and has assumed that firms are observed without error. The literature has also tended to assume that all firms operate precisely the same technology, with all parameters of the production function known exactly. In our context, these assumptions would lead to flawed conclusions. Even though our data have been carefully collected with highly trained enumerators – and although they are often characterized as "state of the art" surveys – measurement error is pronounced, and shocks are quantitatively important.

There are limits to our analysis. As noted in the introduction, we cannot rule out the importance of misallocation in a dynamic sense. The current allocation of land

and labor across farms may be relatively efficient in a static sense, but improved technologies might be well suited to very different allocations. For instance, mechanization and tractor use might increase efficiency in these countries, but it is possible that the current distribution of land might make it unprofitable to use tractors and might thus slow the diffusion of the new technologies.[27] Thus, one could think about a dynamically optimal allocation, which would raise different issues than those we have addressed here.

This paper also suggests that within the literature on agriculture and development, there is a need to pay close attention to heterogeneity in unobservable characteristics of plots. These may be linked to soil and land quality, which vary in quantitatively significant ways at very fine geographic scale. But there may also be a high degree of spatial variation in shadow prices (reflecting, for example, within-farm transport costs). For instance, the distances from one end of a plot to another may create consequential transport and transaction costs for the application of organic fertilizers or for the shadow price of output that must be carried to the household or to market. The importance of heterogeneity has been emphasized in recent work on technology adoption (e.g., Suri 2011), and it is surely important for other issues in agricultural development.

In further work, an interesting area to explore is the trade-off between farm scale and the precision of input application. Because input use is (optimally) calibrated to the average quality of a plot, there is a trade-off between increasing the size of the plot (which reduces the fixed cost per unit output) and the loss of profits that comes from applying inputs more crudely. This trade-off may have some power in explaining the tendency of smallholder agriculture in the developing world to rely so heavily on very small plots, finely tuned in terms of crop choice and input use. Previous explanations of small plot size have tended to focus on risk and diversification, but our analysis suggests that there may also be important efficiency arguments.

[27] We note, however, that tractor use has become relatively widespread in northern Ghana, in spite of the small farm size, and a lively market has merged in tractor services (Cossar, 2017).

8. Conclusions

This paper has examined the importance of misallocation across firms as an explanation for low aggregate productivity in developing countries, using data from agriculture in Ghana, Uganda, and Tanzania. A challenge in this kind of analysis is that misallocation is not the only potential source of dispersion in productivity. Some of the other sources of dispersion are not susceptible to improvement through efficient reallocation. In particular, reallocation will not lead to increases in output if dispersion is primarily an artefact of measurement error. Reallocation will also prove futile to the extent that dispersion results from idiosyncratic shocks that occur after inputs have been (efficiently) applied.

Our paper takes advantage of rich data at the plot level to disentangle the different sources of productivity dispersion. We begin by showing that dispersion in productivity is not simply a feature of the cross-farm data; perhaps surprisingly, within-farm dispersion is quite large. This suggests that differences in farmer quality are not sufficient to account for the patterns of dispersion that we observe in the data.

We proceed to carry out a structural estimation of agricultural production functions for Ghana and Uganda, with a framework that draws on the sequential nature of production decisions. The estimated production functions can be used to assess the potential gains from reallocation. Our finding is that misallocation does indeed affect aggregate agricultural output in these countries, but misallocation accounts for only about 3 percent of total dispersion in Ghana and about 6 percent of total dispersion in Uganda. By our estimates, reallocation can generate substantial gains in aggregate output (not taking into account private or social costs associated with reallocation), but it cannot account for any substantial fraction of cross-country income differences.

Beyond the rather special case of African agriculture, this research points to the need for caution in estimating the impact of misallocation. Not all dispersion in productivity at the firm level reflects misallocation. It is important to for researchers to consider other sources of productivity dispersion, including heterogeneity and measurement error.

References

Ackerberg, Daniel A., Kevin Caves, and Garth Frazer. 2015. Identification properties of recent production function estimators. *Econometrica* 83(6): 2411-2451.

Adamopoulos, Tasso and Diego Restuccia, 2014. The size distribution of farms and international productivity differences. *American Economic Review* 104(6): 1667-1697.

Adamopoulos, Tasso and Diego Restuccia. 2015. Land reform and productivity: A quantitative analysis with micro data. Working Paper 540, University of Toronto, Department of Economics.

Adamopoulos, Tasso, Loren Brandt, Jessica Leight, and Diego Restuccia. 2017. Misallocation, selection and productivity: A quantitative analysis with panel data from China. NBER Working Paper 23039. Cambridge, Mass.: National Bureau of Economic Research.

Asker, John, Allan Collard-Wexler, and Jan De Loecker. 2014. Dynamic inputs and resource (mis) allocation. *Journal of Political Economy*, 122(5): 1013-1063.

Banerjee, Abhijit V. and Benjamin Moll. 2010. Why does misallocation persist? *American Economic Journal: Macroeconomics* 2(1): 189-206.

Barrett, Christopher B., Marc F. Bellemare, and Janet Y. Hou. 2010. Reconsidering conventional explanations of the inverse productivity–size relationship. *World Development* 38(1): 88-97.

Beegle, Kathleen, Calogero Carletto, and Kristen Himelein. 2012a. Reliability of recall in agricultural data. *Journal of Development Economics* 98(1): 34-41.

Beegle, Kathleen, Joachim De Weerdt, Jed Friedman, and John Gibson. 2012b. Methods of household consumption measurement through surveys: Experimental results from Tanzania. *Journal of Development Economics* 98(1): 3-18.

Benjamin, Dwayne. 1992. Household composition, labor markets, and labor demand: testing for separation in agricultural household models. *Econometrica* 60(2): 287-322.

Benjamin, Dwayne. 1995. Can unobserved land quality explain the inverse productivity relationship? *Journal of Development Economics*, 46(1): 51-84.

Bento, Pedro and Diego Restuccia. 2016. Misallocation, establishment size, and productivity. NBER Working Paper 22809. Cambridge, Mass.: National Bureau of Economic Research.

Bevis, Leah and Christopher B. Barrett. 2017. Close to the edge: do behavioral explanations account for the inverse productivity relationship? Mimeo: Cornell University. Paper presented at the 2017 Centre for the Study of African Economies (CSAE) Conference.

Caselli, Francesco. 2005. Accounting for cross-country income differences. In Philippe Aghion & Steven Durlauf (eds.), *Handbook of Economic Growth*, Edition 1, Volume 1, Chapter 9, pp. 679-741. Amsterdam: North-Holland.

Chayanov, A. V. 1966. *The Theory of Peasant Economy*. New York: Irwin. (Originally published in Russian in 1925.)

Chen, Chaoran. 2016. Untitled land, occupational choice, and agricultural productivity. Working Paper 573. University of Toronto, Department of Economics.

Claassen, Roger and Richard E. Just. 2011. Heterogeneity and distributional form of farm-level yields. *American Journal of Agricultural Economics* 93 (1): 144-160.

Collard-Wexler, Allan and Jan De Loecker. 2016. Production function estimation with measurement error in inputs. NBER Working Paper No. w22437. Cambridge, Mass.: National Bureau of Economic Research.

Cossar, Frances. 2017. The role of social capital in the tractor service rental market: A case study from Northern Ghana. Mimeo, Oxford Department of International Development, Oxford, UK.

De Nicola, Francesca, and Xavier Giné. 2014. How accurate are recall data? Evidence from coastal India. *Journal of Development Economics* 106: 52-65.

Deininger, Klaus, Calogero Carletto, Sara Savastano, and James Muwonge. 2012. Can diaries help in improving agricultural production statistics? Evidence from Uganda. *Journal of Development Economics* 98(1): 42-50.

Dillon, Brian and Christopher B. Barrett. 2017. Agricultural factor markets in Sub-Saharan Africa: An updated view with formal tests for market failure. *Food Policy* 67(February): 64-77.

Emran, M. Shahe and Forhad J. Shilpi. 2015. Do land market restrictions hinder structural change in a rural economy? Evidence from Sri Lanka. Policy Research Working Paper Series 7525. Washington, DC: The World Bank.

Foster, Lucia, Cheryl Grim, John Haltiwanger and Zoltan Wolf. 2016. Firm-level dispersion in productivity: Is the devil in the details? *American Economic Review* 106(5):95-98.

García-Santana, Manuel and Josep Pijoan-Mas. 2014. The reservation laws in India and the misallocation of production factors. *Journal of Monetary Economics* 66(C): 193-209.

Gollin, Douglas, David Lagakos, and Michael Waugh. 2014. The agricultural productivity gap. *Quarterly Journal of Economics* 129(2): 939-993.

Gottlieb, Charles and Jan Grobovsek. 2016. Communal land and agricultural productivity. ESE Discussion Papers 270, Edinburgh School of Economics, University of Edinburgh.

Gourlay, Sydney, Talip Kilic, and David Lobell. 2017. Could the debate be over? Errors in farmer-reported production and their implications for inverse scale - productivity relationship in Uganda. Mimeo: Living Standards Measurement Study (LSMS), Survey Unit, Development Data Group, World Bank, Rome, Italy.

Haltiwanger, John. 2016. Firm dynamics and productivity: TFPQ, TFPR, and demand-side factors. *Economía*, 17(1): pp. 3-26.

Hanna, Rema, Sendhil Mullainathan, and Joshua Schwartzstein. 2014. Learning through noticing: Theory and evidence from a field experiment. *Quarterly Journal of Economics* 129(3): 1311-1353.

Heckman, James, and Edward Vytlacil. 1998. Instrumental variables methods for the correlated random coefficient model: Estimating the average rate of return to schooling when the return is correlated with schooling. *Journal of Human Resources* 33(4): 974-987.

Hopenhayn, Hugo A. 2014. Firms, misallocation, and aggregate productivity: A review. *Annual Review of Economics* 6(1): 735-770.

Hsieh, Chang-Tai and Peter Klenow. 2009. Misallocation and manufacturing TFP in China and India. *Quarterly Journal of Economics* 124(4): 1403-1448.

Hurley, Terrance M., Gary L. Malzer, and Bernard Kilian. 2004. Estimating site-specific nitrogen crop response functions. *Agronomy Journal* 96(5): 1331-1343.

Jayachandran , Seema . 2006. Selling labor low: Wage responses to productivity shocks in developing countries. *Journal of Political Economy* 114(3): 538-575.

Kalemli-Ozcan, Sebnem and Bent E. Sorensen. 2012. Misallocation, property rights, and access to finance: Evidence from within and across Africa. NBER Working Paper 18030. Cambridge, Mass.: National Bureau of Economic Research.

Karlan, Dean, Robert Osei, Isaac Osei-Akoto, and Christopher Udry. 2014. Agricultural decisions after relaxing credit and risk constraints. *Quarterly Journal of Economics* 129 (2): 597-652.

Karlan, Dean, Adam Osman, and Jonathan Zinman. 2016. Follow the money not the cash: Comparing methods for identifying consumption and investment responses to a liquidity shock. *Journal of Development Economics* 121(July): 11-23.

Kaur, Supreet. 2016. Nominal wage rigidity in village labor markets. NBER Working Paper No. 20770. Cambridge, Mass.: National Bureau of Economic Research.

Kilic, Talip, Alberto Zezza, Calogero Carletto, and Sara Savastano. 2016. Missing (ness) in action: selectivity bias in GPS-based land area measurements. *World Development* 92(April): 143-157.

LaFave, Daniel, and Duncan Thomas. 2016. Farms, families, and markets: New evidence on completeness of markets in agricultural settings. *Econometrica* 84(5): 1917-1960.

Levinsohn, James, and Amil Petrin. 2003. Estimating production functions using inputs to control for unobservables. *Review of Economic Studies* 70(2): 317-341.

Masten, Matthew A., and Alexander Torgovitsky. 2016. Identification of instrumental variable correlated random coefficients models. *Review of Economics and Statistics* 98(5): 1001-1005.

McMillan, Margaret, Dani Rodrik, and Íñigo Verduzco-Gallo. 2014. “Globalization, structural change, and productivity growth, with an update on Africa.” *World Development* 63(C): 11-32.

Midrigan, Virgiliu and Daniel Yi Xu. 2014. Finance and misallocation: Evidence from plant-level data. *American Economic Review* 104(2): 422-458.

Olley, G. Steven, and Ariel Pakes. 1992. The dynamics of productivity in the telecommunications equipment industry. NBER Working Paper No. w3977. Cambridge, Mass.: National Bureau of Economic Research.

Petrin, Amil, and Jagadeesh Sivadasan. 2013. Estimating lost output from allocative inefficiency, with an application to Chile and firing costs. *Review of Economics and Statistics* 95(1): 286-301.

Petrin, Amil, T. Kirk White, and Jerome P. Reiter. 2011. The impact of plant-level resource reallocations and technical progress on US macroeconomic growth. *Review of Economic Dynamics* 14(1): 3-26.

Restuccia, Diego and Richard Rogerson. 2008. Policy distortions and aggregate productivity with heterogeneous establishments. *Review of Economic Dynamics* 11(4): 707–20.

Restuccia, Diego and Richard Rogerson. 2013. Misallocation and productivity. *Review of Economic Dynamics* 16(1): 1-10.

Restuccia, Diego and Raul Santaeulalia-Llopis. 2017. Land misallocation and productivity. NBER Working Paper 23128. Cambridge, Mass.: National Bureau of Economic Research.

Restuccia, Diego, Dennis Tao Yang, and Xiadong Zhu. 2008. Agriculture and aggregate productivity: A quantitative cross-country analysis. *Journal of Monetary Economics* 55(2): 234-50.

Schimmelpfennig, David. 2016. *Farm Profits and Adoption of Precision Agriculture*. No. 249773. United States Department of Agriculture, Economic Research Service.

Schultz, T.W. 1964. *Transforming Traditional Agriculture*. New Haven: Yale University Press.

Stoorvogel, Jetse J., Lammert Kooistra, and Johan Bouma. 2015. Managing soil variability at different spatial scales as a basis for precision agriculture. In *Soil-Specific Farming: Precision Agriculture*, eds. R. Lal and BA Stewart: 37-72.

Suri, Tavneet. 2011. Selection and comparative advantage in technology adoption. *Econometrica* 79(1): 159-209.

Syverson, Chad. 2004. Product substitutability and productivity dispersion. *Review of Economics and Statistics* 86(2): 534-50.

Syverson, Chad. 2011. What determines productivity? *Journal of Economic Literature* 49(2): 326-365.

Tittonell, P.D., K.D. Shepherd, B. Vanlauwe, and K.E. Giller. 2008. Unravelling the effects of soil and crop management on maize productivity in smallholder agricultural systems of western Kenya—An application of classification and regression tree analysis. *Agriculture, Ecosystems & Environment* 123(1-3): 137-150.

Tittonell, Pablo, Bernard Vanlauwe, P. A. Leffelaar, Keith D. Shepherd, and Ken E. Giller. 2005. Exploring diversity in soil fertility management of smallholder farms in western Kenya: II. Within-farm variability in resource allocation, nutrient flows and soil fertility status. *Agriculture, Ecosystems and Environment* 110(3): 166-184.

Tittonell, Pablo, Bernard Vanlauwe, N. De Ridder, and Ken E. Giller. 2007. Heterogeneity of crop productivity and resource use efficiency within smallholder Kenyan farms: Soil fertility gradients or management intensity gradients? *Agricultural Systems* 94(2): 376-390.

Townsend, Robert M. 1994. Risk and insurance in village India. *Econometrica* 62(3): 539-591.

Udry, Christopher. 1994a. Gender, agricultural production, and the theory of the household. *Journal of Political Economy* 104(5): 1010-1046.

Udry, Christopher. 1994b. Risk and insurance in a rural credit market: An empirical investigation in northern Nigeria. *The Review of Economic Studies*, 61(3): 495-526.

Udry, Christopher. 1996. Efficiency in market structure: testing for profit maximization in African Agriculture. Mimeo: Northwestern University.

Vanlauwe, Bernard, P. Tittonell, and J. Mukalama. 2006. Within-farm soil fertility gradients affect response of maize to fertiliser application in western Kenya. *Nutrient Cycling in Agroecosystems* 76(2): 171-182.

Vanlauwe, B., K. Descheemaeker, K. E. Giller, J. Huising, R. Merckx, G. Nziguheba, J. Wendt, and S. Zingore. 2015. Integrated soil fertility management in sub-Saharan Africa: unravelling local adaptation. *Soil* 1(1): 491-508.