The inverse productivity size relationship: can it be explained by systematic measurement error in self-reported production?

Sam Desiere

Invited paper presented at the 5th International Conference of the African Association of Agricultural Economists, September 23-26, 2016, Addis Ababa, Ethiopia

Submission to the 5th International Conference of the African Association of Agricultural Economists (ICAAAE)
23-26 September 2016
Addis Ababa, Ethiopia

The inverse productivity size relationship: can it be explained by systematic measurement error in self-reported production?

Sam Desiere
Department of Agricultural Economics, Ghent University, Belgium
sam.desiere@ugent.be

Abstract
This paper revisits the decades-old puzzle of the inverse productivity plot-size relationship (IR), which states that land productivity decrease with increasing plot size in developing countries. While most empirical studies about the IR define yields as self-reported production divided by plot size, this paper complements this approach with an alternative, objective method to estimate yields: crop cuts. Using crop cuts as proxy for yields, the IR in Ethiopia disappears, while the relationship is strong when yields are based on self-reported production. The inverse relationship is even reversed as there exists a weak, positive correlation between plot size and crop cuts. This implies that farmers systematically over report production on small plots and underreport it on larger ones. Our findings suggest that the IR is an artifact of systematic measurement error in self-reported production.

1. Introduction

The inverse relation (IR) between farm size and land productivity has puzzled agricultural economists for over half a century. As first noted by Chayanov (1926) in Russia and rediscovered by Sen (1962) in India, it states that production per hectare decreases with increasing farm and, even, plot size. These relationships have been observed in many developing countries over several decades and are often considered a stylized fact (Larson, Otsuka, Matsumoto, & Kilic, 2013).

The policy implications of the inverse productivity-size relationship are disturbing and counter-intuitive for many agricultural economists, which explains the vast literature on the topic. An oft-emphasized implication concerns land redistribution. Redistributing land from large-scale to small-scale farmers will not only improve equality, but may also increase production since small-scale farmers are more efficient than large-scale farmers. In a similar vein, the IR suggests that policy makers should focus on small-scale farmers to generate productivity growth, rather than on promoting large-scale agriculture. Although agricultural economists are well aware of the IR, few would recommend land redistribution for efficiency reasons and some even argue that large-scale farming is a more promising pathway for rural development in developing countries (Collier & Dercon, 2013).

The implication of the IR for land consolidation programs are also important, although less frequently emphasized (Ali & Deininger, 2014; Blarel, Hazell, Place, & Quiggin, 1992). In pursuit of productivity gains and economies of scales, several countries have implemented land consolidation programs to halt land fragmentation and create larger consolidated fields. The IR predicts that this strategy will fail since small plots are more productive than large plots. Again, it is not unusual that agricultural economists and policymakers favor land consolidation in regions where the agricultural landscape is highly fragmented (Pašakarnis & Maliene, 2010).

The IR is even at odds with economic theory. One of the core principles of micro-economy is that factor productivity should be equal across farms and plots. In the absence of factor productivity equality, active land markets should assure that households with higher marginal productivity acquire land from households with a lower marginal productivity until the marginal returns on land are equal across farms. Within a single household, marginal productivity should also be equalized across plots since rational households allocate labor and inputs efficiently across plot (Barrett, Bellemare, & Hou, 2010).

Because of the controversial policy implications of the IR as well as its inconsistency with economic theory, a vast theoretical and empirical literature has emerged offering and testing explanations for the IR that are consistent with conventional wisdom. These explanations can be classified in three sets. The most common explanation for the inverse productivity farm-size relationship is related to missing or imperfect land and labor markets (Carter & Wiebe, 1990), credit and insurance markets (Barrett, 1996) and imperfections in labor supervision (Eswaran & Kotwal, 1986; Feder, 1985). The consequence of imperfect markets is that small-scale farmers apply more than the optimal amount of inputs on their plots since their outside options are limited. Yet, while missing markets can explain differences in land productivity between households, they cannot explain productivity differences across plots within a single household (Assunção & Braido, 2007). The second explanation, differences in soil quality between small and large plots, can explain both the inverse productivity plot-size and farm-size relationship (Benjamin, 1995; Bhalla & Roy, 1988). Soil quality explains the IR if smaller farms (plots) have on average more fertile soils than larger farms (plots). Since it is notoriously difficult to measure soil quality, most empirical work omits this variable altogether or uses self-reported measures. Using excellent, objective soil quality data, Barrett et al. (2010) showed, however, that soil quality contributes only marginally to explaining the IR in Madagascar. The final explanation for the IR is related to measurement error in self-reported land area, which was first suggested by Lamb (2003). Recent work showed that measurement error in land alone is unlikely to explain the inverse relationship. Replacing self-reported land size with GPS measurement weakened the IR somewhat in several African countries (Carletto, Gourlay, & Winters, 2015), but the relationship never disappeared and even strengthened in Uganda (Carletto, Savastano, & Zezza, 2013).

This paper offers and tests a new explanation for the inverse productivity plot-size relationship: measurement error in self-reported production. If farmers systematically over report production on small plots and underreport it on larger plots, this would generate a spurious inverse productivity plot-size relationship. As far as we could assess[1], recent studies on the IR defined yields as self-reported production divided by plot size (e.g. (Ali & Deininger, 2015; Carletto, Gourlay, et al., 2015; Carletto et al., 2013; Larson et al., 2013)). Systematic measurement error in production may therefore affect the estimation of the IR in several studies.

One reason why systematic measurement error in self-reported production has not yet been considered as an explanation for the IR are the demanding data requirements to test it. Systematic measurement error can only be observed if one has a second, independent measurement method. In this study, we draw on two waves of a nationally representative dataset from Ethiopia which asked farmers to report production for all plots, but which also implemented crop cuts on a limited set of plots. Crop cuts are an objective method to measure yields. Yields are estimated by randomly sampling a small subplot (often a 4m x 4m square) within the plot, delimiting it and cutting and weighing the harvest of this subplot. Yields are then defined as the harvest in the subplot divided by the area of the subplot. Crop cuts are sometimes considered the gold standard of yield measurement (Fermont & Benson, 2011).

Using crop cuts as proxy for yields, the inverse productivity plot-size relationship disappears in our data, while this relationship is strong if yields are based on self-reported production. After carefully establishing the robustness of these findings and evaluating alternative explanations, most notably differences in labor input between plots, systematic measurement error in self-reported production remains the only plausible explanation. The key assumption for this explanation to hold is that measurement errors in crop cuts occur randomly, while measurement error in self-reported production can be systematic and may be correlated with plot size. This assumption holds by definition if crop cuts are the gold standard of yield measurement. Even if they are not, it appears unlikely that measurement errors in crop cuts vary with plot size. Bias in self-reported variables, on the other hand, has been shown to correlate with plot and household characteristics. For instance, households systematically overestimate the area of small plots and underestimate the area of larger plots (Carletto,

[1] Most studies do not discuss how they measured yields. We contacted the authors of the several recent studies on the IR cited above. They confirmed that yields were based on self-reported production.

Gourlay, et al., 2015), while larger households underreport consumption relative to smaller households (Gibson & Kim, 2007).

In addition to advancing the literature on the IR, this paper contributes to the small literature on systematic measurement errors in household surveys and their implications for statistical inference. This paper adds to a set of papers that have demonstrated that not all measurement error is white noise, which may induce spurious correlations in the data (Aguiar & Bils, 2015; Beegle, De Weerdt, Friedman, & Gibson, 2012; Carletto, Gourlay, et al., 2015; Gibson & Kim, 2007). As is well-known, but perhaps insufficiently emphasized, this illustrates that the 'Iron Law of Econometrics' – as Hausman famously called the observation that estimated coefficients in a regression usually underestimate the true magnitude – critically depends on the assumption that measurement error occurs randomly (Hausman, 2001). If measurement error is systematic, spurious correlations may be observed in the data.

The remainder of this paper is structured as follows. In the next section, we discuss the data and provide some descriptive statistics. Next, we outline our estimation strategy and discuss the conditions under which systematic bias in self-reported production can generate the inverse productivity plot-size relationship. In the result section, we show that the inverse productivity plot-size relationship exists when yields are based on self-reported production, but disappears when yields are estimated with crop cuts. Based on these findings, the sign and magnitude of the measurement error in self-reported production as function of plot size is discussed. Section 5 concludes.

2. Data and descriptive statistics

We use data from two waves of the nationally representative Ethiopian Socioeconomic Survey (ESS). This survey is an ongoing project to collect high-quality panel data in Ethiopia. It is implemented by the Central Statistical Agency of Ethiopia in close collaboration with the LSMS-ISA[2] team of the World Bank, which has a long history of producing high-quality data. All data and relevant documentation is publicly available.

The first wave was administered in 2011/2012 to 3969 rural households, while the second wave was administered in 2013/2014 to 5262 household – 3776 panel households and 1486 new, mainly urban, households. We only included households in the dataset that were interviewed in both waves. The time dimension in the data was not exploited and all the observations were simply pooled. The survey collected standard information on household characteristics, consumption, living conditions and health.

The unique feature of the survey is its focus on agriculture. To gather detailed and accurate agricultural data at plot level, households were visited three times during the agricultural year. The first visit occurred in September-October to collect data on planting activities. During this visit, the area of most plots was measured with GPS. The second visit occurred in November and implemented the livestock module. The final visit took place from January to April and collected data on agricultural production. This visit also included the household questionnaire.

[2] LSMS-ISA: Living Standards Measurement Study – Integrated Surveys on Agriculture.

In this paper, we exploit that yields were measured with two different methods: crop cuts and self-reported production by plot in combination with land measurement with GPS3. In addition, we also used the detailed information on labor input during planting (first visit) and the harvest (third visit).

Crop cuts were implemented in both waves of the survey for 23 major crops. Five plots per crop were randomly selected from a list of all plots cultivated by the sampled households within an enumeration area. In most cases, the plot selected for crop cutting was monocropped. Only if the crop was cultivated on less than five plots in pure stand within an enumeration area, were crop cuts also implemented on intercropped plots. Once a plot was sampled from this list, a rectangular subplot within the plot was randomly sampled. In the first wave, this subplot was $4m^2$ (2m x 2m), while it was $16m^2$ (4m x 4m) in the second wave. Within the subplot, the crop was harvested by a trained enumerator and weighed. If logistical constraints allowed, crop cutting occurred simultaneously with the harvest of the main crop by the farmer. In both waves, nearly 40% of the crop cuts were executed in November and more than 90% were executed between October and December. Both fresh weights and dry weights were recorded. The correlation between both measures was over 0.95. We used the dry weights to calculate yields. 2975 and 3532 crop cuts were taken in wave 1 and 2, respectively. We discarded, however, those crop cuts for which we had fewer than 150 observations per crop. Finally, 5920 crop cuts remained in the dataset, providing yield estimates for 19 different crops.

Farmers reported the harvest per crop and plot during the third visit. In wave 1, most visits occurred in January (57%), while in wave 2 most visits took place in February (78%). Note that the recall period ranges up to 5 months since most crops during the Meher season are harvested from September till February (Taffesse, Dorosh, & Asrat, 2011). On average, production was reported 77 and 98 days after the implementation of the crop cuts in wave 1 and 2, respectively.

Production was reported on 3683 and 23,638 plots in wave 1 and 2, respectively. Besides production data, detailed data on labor input during planting and the harvest at plot level was also collected.

As we focus on the inverse productivity plot-size relationship, the unit of analyses is the plot. It is therefore important to clearly define a 'plot', as this definition is context specific. In the Ethiopian household survey, enumerators first defined parcels, which are units of land that are owned by a single household and surrounded by land owned by another household or demarcated by natural boundaries (forest, water, road). Within a parcel, the plots were identified. Plots were clearly demarcated by hedges or paths. In most plots (80%) a single crop was cultivated. All data (land area, production, inputs and crop cuts) was collected at plot level, with the exception of irrigation, soil quality and the possession of a land certificate, which formalizes land tenure, which were reported at parcel level. Since parcels are already small pieces of land, plots are even tinier. Mean and median plot size is 0.13ha and 0.064ha, respectively and more than 95% of the plots were smaller than 0.5ha. Production at plot level was valued at median self-reported prices in order to make output from different crops comparable.

In order to examine if crop cuts were indeed randomly implemented across households and plots, we compare households with at least one plot selected for crop cutting versus those without a single plot selected for crop cutting and compare plots with and without crop cuts. Household characteristics are very similar between both groups (table 1). The only important difference is that households selected for crop cuts owned slightly more land (1.30 ha versus 1.19 ha). This is in line with expectations since households with more land are also more likely to cultivate at least one plot suited for crop cutting. In terms of plot characteristics, differences between plots selected for crop cuts and the other plots are also due to the sampling design (table 2). Plots with crop cuts are larger, further away from the dwelling of the household and more likely to be planted with a single crop. Differences in terms of fertilizer application (both organic and inorganic), irrigation, labor input and soil quality are small. In sum, crop cuts occurred randomly across households and plots.

Table 1: Household characteristic for households with and without plots with crop cuts

	No crop cuts	Crop cuts
Landholdings (ha)	1.19	1.30
	(0.052)	(0.076)
Applied chemical fertilizer (%)	0.47	0.491
	(0.019)	(0.017)
Asset index	0.16	0.16
	(0.0016)	(0.0024)
Household size	5.79	5.75
	(0.052)	(0.076)
Age household head	46.36	45.82
	(0.33)	(0.50)
Household head can read and write	38	38
	(0.010)	(0.016)
Female headed household (%)	20	17
	(0.0086)	(0.012)
N	2069/2125	905/918

Standard errors in parentheses. Number of observations differs by variable due to missing variables

Table 2: Plot characteristic for plot selected and not selected for crop cutting

	Plots without crop cuts	Plots with crop cuts
Plot size (m^2)	1227	1501
	(10.9)	(21.6)
Log of distance	0.383	0.519
	(0.003)	(0.007)
Plot slope (%)	14.6	14.3
	(0.084)	(0.159)
Plot elevation	1931	1968
	(3.24)	(6.38)
Plot potential wetness index	12.6	12.5
	(0.013)	(0.026)
Land title (certificate)	0.483	0.502
	(0.003)	(0.007)
Pure stand	0.572	0.867
	(0.003)	(0.004)
Applied manure (% plots)	33.1	20.7
	(0.003)	(0.005)
Applied compost (% plots)	5.2	5.8
	(0.002)	(0.003)
Applied organic fertilizer (% plots)	1.8	2.0
	(0.001)	(0.002)
Irrigation (% plots)	4.4	2.3
	(0.001)	(0.002)
Self-reported soil quality (only wave 2)	1.748	1.890
	(0.005)	(0.012)
Fertilizer (kg/ha)	44.4	46.2
	(0.752)	(1.384)
Distance to dwelling (km)	0.685	0.953
	(0.007)	(0.016)
Family labor planting (days/ha)	16.9	16.1
	(0.173)	(0.313)
Hired labor planting (days/ha)	7.5	7.3
	(0.299)	(0.603)
Exchange labor planting (days/ha)	15.2	15.1
	(0.395)	(0.670)
Family labor harvesting (days/ha)	12.2	8.1
	(0.121)	(0.166)
Hired labor activities (days/ha)	4.8	4.16
	(0.218)	(0.337)
Exchange labor activities (days/ha)	15.34	15.08
	(0.374)	(0.642)
N (max	21401	5920
N (min)	20046	3296

Standard errors in parentheses. Number of observations differs by variable due to missing variables. Soil quality was only reported in wave 2.

In order to explore if an inverse productivity plot-size relationship exists in our data, we examined maize yields in wave 2 by quartile of plot size and measurement method (figure 1). Mean and median yields based on self-reported production clearly decrease with increasing plot size. Median maize yields are over 2000 kg/ha on the smallest plots, but decrease to 1000 kg/ha on the largest plots.

The inverse productivity plot size relationship disappears, however, when yields are estimated with crop cuts. Crop cuts of maize even increase with plot size. When comparing crop cuts with yields based on self-reporting, we note that crop cuts are substantially lower than self-reported yields on the smallest plots, but higher on the largest plots. This suggests that self-reported production is systematically overestimated on small plots and underestimated on larger plots. Given the few important differences between households and plots with or without crop cuts, this suggests that the inverse productivity size relationship can be attributed to systematic measurement error in self-reported production. This will be examined in more detail in the remainder of this paper.

Figure 1: Mean (left panel) and median (right panel) maize yields (kg/ha) in wave 2 by measurement method in function of plot size

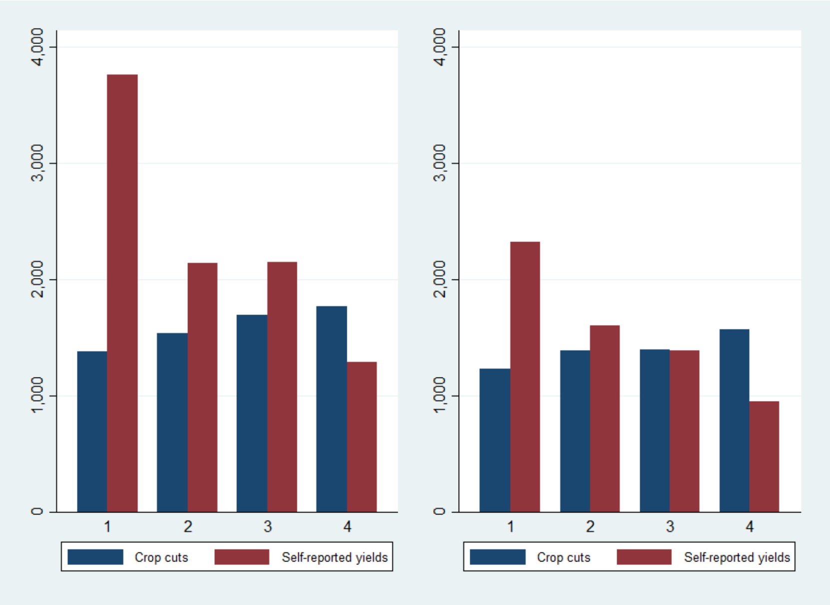

Sample restricted to plots with both crop cuts and self-reported production. 62, 134, 138 and 134 observations in quartile 1 to 4, respectively. The quartiles have an average plot size of 150m², 425m², 1128m² and 3645m², respectively.

3. Methodology

The standard approach to estimate the inverse productivity size relationship is straightforward: yields are regressed on plot size (A) and a set of plot and household characteristics (Assunção & Braido, 2007; Barrett et al., 2010). Hence, the following equation is typically estimated:

$$\log(yield_{ijr}) = \alpha\log(A_{ij}) + \beta X_i + \delta Z_j + \mu L_j + v_r + \epsilon_{ij} \qquad (1)$$

Where i, j and r represent the household, the plot and the enumeration area, respectively. The coefficient α is the parameter of interest which defines the strength of the inverse relation, X_i is a vector of household characteristics, Z_j represents plot characteristics, L_j is labor input at plot level and v_r are enumeration area fixed effects.

This equation will be estimated twice using a different proxy for yields: self-reported production divided by plot size or crop cuts. Both approaches for yield measurement are imperfect for different reasons and measure yields with error (Fermont & Benson, 2011). The properties of these errors determine if measurement error in yields can induce a spurious correlation between plot size and land productivity.

One has to make a distinction between two types of measurement error: systematic measurement error or bias and random measurement error (Boumans, 2015). Bias reduces the accuracy of the measurement, while random measurement error reduces the precision of the measurement. Ideally, measurement instruments are both accurate and precise. One intuitive approach to formalize this distinction is called the Mean Squared Error (MSE), which is defined as the square of the expected difference between the outcome of the measurement and the 'true' (unobserved) value of the concept that is being measured (De Groote & Traoré, 2005):

$$MSE(yield) = E(yield - yield^*)^2 = \left(yield^* - E(yield)\right)^2 + V(yield) \qquad (2)$$

Where $yield$ is the outcome of the measurement using self-reporting or crop cuts, and $yield^*$ represents true, unobserved yields.

The first term in equation 2 shows the bias in the measurement, that is, it shows the systematic difference between the true yields and the expected value of the measurement. The second term is the variation in the measurement, which corresponds to the inverse of the measurement precision. While measurement precision increases with sample size (e.g. repeated measurement), the bias is independent of sample size and an integral part of the measurement instrument.

Both the bias and the precision of the measurement can be correlated with plot size. However, only a correlation between the bias and plot size can induce a spurious inverse productivity size relationship (Carroll, Ruppert, Stefanski, & Crainiceanu, 2012). To see this, assume that the bias decrease with plot size. For instance, the measurement always overestimate true yields on small plots and correctly estimates true yields on larger plots. This creates a negative correlation between estimated yields and plot size. More formally, a correlation between bias in the measurement and plot size introduces a correlation between the error term in equation 1, ε, and plot size, which biases the parameter of interest, α.

The variance in yield measurement, the second term in equation 2, may also be correlated with plot size. For instance, yields might be less precisely estimated on larger plots. This creates heteroscedasticity in the error term in equation 1. Heteroscedasticity does not lead to biased estimates, but does bias standard errors. One can correct for heteroscedasticity by using robust estimation techniques.

In sum, a negative correlation between bias in yield measurement (first term in equation 2) and plot size cause a spurious inverse productivity plot size relation, while a correlation between the variance of the measurement (second term in equation 2) and plot size causes heteroscedastic errors. Crop cuts and self-reported production are both biased and imprecise estimates of unobserved yields. The question is whether the bias correlates with plot size.

Measurement error enters self-reported production because respondents may not accurately recall how much they harvested or because they systematically under or overestimate production. Importantly, previous research has shown that errors in self-reported variables can systematically be correlated with crop type and household or plot characteristics. Carletto, Gourlay, et al. (2015), for instance, show that households systematically overestimate plot size on small plots and underestimate it on larger ones. Gibson and Kim (2007) analyses suggest that larger households systematically underestimate consumption relative to smaller households, while Beegle et al. (2012) illustrate how different approaches of consumption measurement can substantially affect poverty estimates. Several authors have argued that crops harvested on a daily basis such as cassava are less precisely recalled than crops that are harvested at one moment in time (Carletto, Jolliffe, & Banerjee, 2015; Jerven, 2013). These correlations introduces systematic bias, which invalidate statistical inference. It cannot be excluded that similar processes cause a negative correlation between bias in self-reported production and plot size[4].

In contrast to bias in self-reported production, bias in crop cuts is unlikely to be correlated with household or plot characteristics. Enumerators always followed the same guidelines when implementing the crop cuts, which are completely independent of plot size. Yet, crop cuts introduce a type of error that is absent for self-reported production: sampling errors. Sampling errors affect the precision of the measurement, but do not cause bias. Crop cuts are taken from a randomly selected subplot within the plot. Since yields are likely to vary within a plot, this introduces sampling error in the yield estimate. Sampling errors introduces a negative correlation between the precision of the measurement and plot size. Consider, for instance, a plot that is exactly equal to the subplot selected for crop cutting (i.e. 16m² in wave 2). In this case, yields are perfectly measured because the whole plot is harvested. Consider next a plot of 1ha (10,000 m²). A crop cut of 16m² within a plot of 1ha is unlikely to give a precise estimate of the yields. Hence, the precision of yield measurement with crop cuts decreases with plot size. The more variation in yields within a plot or the larger the plot, the less precise the estimation of yields. Yet, a lack of precision in the measurement can never cause a spurious correlation between plot size and land productivity since it only increases the standard errors of the estimated coefficients in equation 1.

So far, we implicitly assumed that crop cuts and yields based on self-reported production capture the same underlying concept. In other words, the 'true', unobserved yields, represented by $yield^*$ in equation 2, are assumed to be the same for both methods. This is not necessarily true. Crop cuts measure the maximum potential yield, while self-reported production measures the actual harvest. In theory, crop cuts are an upper bound of yields, while self-reported production gives a more accurate estimate of the realized yields since it takes harvest losses into account. This conceptual difference challenges our empirical analysis. If the IR holds for self-reported proxies of land productivity, but disappears for crop cuts, one may argue that smaller plots are simply more completely harvested. To control for this possibility, we include labor input during planting and harvest at plot level in the regressions. In addition, we estimate the correlation between plot size and labor input in order to examine if small plots are indeed more completely harvested than larger ones.

The previous discussion about biased measurement of yields showed that the inverse productivity size relationship will be observed for yields based on self-reported production and disappear for crop cuts if farmers overestimate production on small plots relative to larger ones. This hypothesis can be tested directly by estimating the following equation:

$$\log\left(yield_j^{SR}/yield_j^{CC}\right) = C + \delta\log(A_j) + \beta L_j + \varepsilon_j \qquad (3)$$

[4] Precision of yield measurement based on self-reported production may be positively correlated with plot size. If self-reported production is proportionally less over or underestimated on small plots than large ones, the correlation between plot size and the precision of the measurement is negative. For instance, if 'true' yields are 1000kg/ha, production on a plot of 1ha equals 1000kg, while it equals 10kg on a plot of 100m². If a farmers over or under estimate production by 1kg on the small plot and by 50kg on the larger plot, yields are over or underestimated by 10% on the small plot and by 5% on the larger plot.

Where $yield_{ij}^{SR}/yield_{ij}^{CC}$ is the ratio of yields based on self-reported production to yields based on crop cuts. If δ is negative, farmers indeed systematically over report production on small plots relative to larger plots. Labor inputs, L_i, is also included in the regression since self-reported yields may be higher than crop cuts on plots that are more completely harvested.

Based on equation 3, one can then derive the sign and magnitude of measurement error in self-reported production. Self-reported yields equal the sum of the 'true', but unobserved production, X^*, and an error term, μ, divided by plot area, A. True, but unobserved yields (yield*) equal true production divided by plot size. We are interested to estimate μ, the absolute error in self-reported production, in function of plot size.

Crop cuts are unlikely to be exactly equal to the true yields. Rather, we assume that the ratio of crop cuts to 'true' yields equals an unknown constant, v. Previous research has shown that crop cuts are likely to overestimate 'true' yields, and we thus believe that $v > 1$ (Fermont & Benson, 2011). Based on these assumptions and by manipulating equation 3, we can now derive measurement error in self-reported production:

$$\frac{\mu}{A_j * yield^*} = \left(v\gamma A_j^{\delta} - 1\right) \quad \text{with } \gamma = \exp(C); \quad -1 < \delta < 0 \; (4)$$

The term $\frac{\mu}{A_j * yield^*}$ represents relative measurement error in self-reported production, that is, the ratio of the absolute measurement error to the 'true' production. Since the right-hand side of equation 4 can be positive or negative, self-reported production can overestimate true production ($\mu > 0$) or underestimate it ($\mu < 0$), depending on the values of the parameters δ, v, γ. Production is under estimated on plots smaller than $\frac{\log(v\gamma)}{\delta}$, and overestimated on plots larger than this critical value. This thresholds depends on the difference in the strength of the IR estimated with self-reported yields and crop cuts. However, since the parameter v is unknown, we cannot determine this threshold exactly without additional assumptions about v. Recall that v is the ratio of yields based on crop cuts to 'true' yields. If one is willing to accept that crop cuts are the gold standard of yield measurement, v exactly equals 1. If crop cuts always overestimate 'true' yields, v is larger than 1. In order to get a rough estimate of the magnitude of relative measurement error as well as the threshold below which self-reported production is overestimated, we calculate relative measurement error and the threshold for several values of v.

4. RESULTS

Table 3 shows the results of the estimation of equation 1 using yields based on self-reported production (column 1-4) and yields based on crop cuts as the dependent variable (column 5-6). In all specifications robust t-statistics are reported to account for heteroscedasticity. We first discuss the results for self-reported yields and firmly establish that the inverse relationship holds in the data. Next, we show that the inverse productivity plot size relationship disappears when yields are measured with crop cuts. We then examine if differences in labor input between small and large plots or systematic bias in self-reported production explain the findings.

The classic approach to estimate the inverse productivity plot-size relationship includes household fixed effects and plot characteristics (table 3, column 1). This shows that the inverse relationship holds in Ethiopia and is stronger than in many other Sub-Saharan African countries (Larson et al., 2013): yields decrease by 40% if plot size doubles. The second specification replaces the household fixed effects by enumeration area fixed effects and includes a set of household characteristics. This does not affect the IR. In other words, the results remain similar when examining differences in land productivity between large and small plots within a household or between households within a same enumeration area. This is important because household fixed effects are replaced by enumeration area fixed effects when using crop cuts as proxy for land productivity. Restricting the sample to plots for which we have both self-reported production and crop cuts (columns 3 and 4), reduces the sample size from over 25,000 observations to slightly over 5000 observations. In both specifications, the IR remains highly significant. Specification 4 includes labor at plot level as additional control variable, which substantially weakens the IR. Without including labor, yields decrease by 30% if plot size doubles, while yields only decrease by 16% if differences in labor input between small and large plots are taken into account. This important point will be discussed in more detail below.

In addition to plot size and labor, there are several plot and household characteristics that are significantly correlated with yields based on self-reported production. Of the 11 plot characteristics included in the regressions, 4 are significantly correlated with yields. Irrigation and the application of fertilizer as well as the distance to the dwelling increase yields. Most plots (30%) are located in the immediate vicinity of the dwelling. It may be that these plots are overexploited and have therefore lower yields. In wave 2, yields are higher on plots with mixed cropping systems, perhaps because of complementary between crops. In wave 1, yields are by design always underestimated on mixed plots. With regards to household characteristics, only the asset index is significantly correlated with yields in most specifications. The asset index is a simple, standardized count of the number of assets owned by the household. Thirty-five assets are considered including agricultural equipment such as ploughs, sickles and axes as well as durable consumption goods such as radios and mobile phones. It is a proxy for household's wealth. Perhaps unsurprisingly, wealthier households tend to have higher yields. This may be because they own more agricultural equipment, are better connected to extension services or because they are less risk averse than poorer households. In some specifications, female headed households have lower yields, while poorly educated household have higher yields.

The results of the estimation of the IR using crop cuts as dependent variable are also shown in table 3 (column 5-6). Specification 5 is the counterpart of specification 3, while specification 6 includes labor input and can be compared directly to specification 4. In both specifications the IR disappears. We even observe a positive and highly significant correlation between plot size and yields. Several other variables are also correlated with crop cuts. As in the previous specifications, chemical fertilizers and the asset index increase yields, although the latter variable is only marginally significant. Labor input also correlates positively with crop cuts, but the correlations are weaker than between labor and yields based on self-reporting.

Table 3: The inverse plot-size relationship for yields based on self-reporting and crop cuts

	Self-reported measurement				Crop cuts	
	(1)	(2)	(3)	(4)	(5)	(6)
Log of plot size (m²)	**-0.397*****	**-0.396*****	**-0.303*****	**-0.161*****	**0.104*****	**0.149*****
	(-55.84)	(-36.11)	(-12.65)	(-6.22)	(8.63)	(9.36)
Wave	-0.103***	-0.0978*	-0.101*	-0.0966*	-0.427***	-0.403***
	(-2.98)	(-1.78)	(-1.87)	(-1.69)	(-9.45)	(-8.60)
Plot characteristics						
Log of distance field to dwelling	0.133***	0.110***	0.150***	0.0925**	-0.0188	-0.0285
	(4.56)	(4.26)	(3.27)	(2.20)	(-0.65)	(-1.01)
Plot Slope (percent)	-0.00238*	0.000461	0.00202	0.00132	-0.000340	-0.000522
	(-1.86)	(0.38)	(0.88)	(0.60)	(-0.24)	(-0.38)
Plot Elevation (m)	-0.000246**	-0.000198*	-0.000262	-0.000148	-0.000328***	-0.000248**
	(-2.27)	(-1.83)	(-1.64)	(-1.04)	(-3.08)	(-2.34)
Plot Potential wetness Index	-0.000669	-0.000929	-0.00625	-0.00684	-0.00882	-0.00952
	(-0.12)	(-0.19)	(-0.61)	(-0.77)	(-1.09)	(-1.25)
Household has land title (no=1)	0.0105	0.0190	0.0193	0.0429	0.0436*	0.0532**
	(0.28)	(0.74)	(0.46)	(1.10)	(1.69)	(2.09)
Crop in pure stand (no=1)	-0.304***	-0.325***	-0.308***	-0.313***	-0.0871	-0.102
	(-4.35)	(-3.34)	(-3.39)	(-3.55)	(-1.09)	(-1.28)
Pure stand (no=1)*Wave 2	0.619***	0.619***	0.457***	0.465***	0.0445	0.0602
	(8.52)	(5.90)	(3.44)	(3.70)	(0.39)	(0.54)
Manure applied (no=1)	0.0456*	0.0191	-0.0129	0.0191	-0.0677*	-0.0573
	(1.82)	(0.68)	(-0.26)	(0.41)	(-1.89)	(-1.60)
Compost applied (no=1)	0.0420	0.0196	-0.116	-0.0952	-0.0455	-0.0223
	(0.94)	(0.48)	(-1.51)	(-1.30)	(-0.68)	(-0.35)
Organic fertilizer (no=1)	-0.0266	0.0663	-0.0841	-0.0799	-0.254**	-0.265**
	(-0.33)	(0.81)	(-0.56)	(-0.58)	(-2.27)	(-2.38)
Field irrigated (no=1)	-0.272***	-0.202**	-0.290***	-0.265**	0.0340	0.0520
	(-4.08)	(-2.15)	(-2.78)	(-2.59)	(0.42)	(0.63)
Log of fertilizer (kg/ha)	0.0838***	0.0779***	0.0840***	0.0496***	0.0421***	0.0272***
	(17.74)	(12.12)	(9.94)	(6.12)	(6.21)	(3.98)
Household characteristics						
Asset index		0.682***	0.591***	0.514**	0.319*	0.345*
		(4.22)	(2.93)	(2.57)	(1.72)	(1.66)
Female headed household (no=1)		0.0801***	0.0533	0.0561	0.00490	0.0151
		(3.12)	(1.02)	(1.16)	(0.14)	(0.45)
Age household head		0.000379	-0.00409	-0.00942	-0.00208	-0.00256
		(0.11)	(-0.65)	(-1.57)	(-0.42)	(-0.51)
Age²		0.00000278	0.0000461	0.0000909	0.0000332	0.0000357
		(0.08)	(0.77)	(1.59)	(0.67)	(0.71)
Household head can read and write (no=1)		0.0503**	0.00557	0.00674	-0.0270	-0.0297
		(2.40)	(0.15)	(0.20)	(-0.95)	(-1.11)
Labor input (logs)						
Family labor planting (days/ha)				0.123***		0.0718***
				(6.06)		(4.79)
Hired labor planting (birr/ha)				0.0365***		0.0210*
				(2.64)		(1.89)
Exchange labor planting (days/ha)				0.00427		0.0158**
				(0.43)		(2.16)
Family labor harvest (days/ha)				0.292***		0.0591***
				(12.23)		(3.31)
Hired labor harvest (birr/ha)				0.0918***		0.0158
				(5.74)		(1.39)
Exchange labor (days/ha)				0.105***		0.0353***
				(9.94)		(5.34)
Constant	12.64***	11.93***	12.34***	10.33***	10.13***	9.278***
	(36.59)	(28.08)	(19.90)	(18.36)	(23.98)	(21.31)
Household fixed effects	Yes	No	No	No	No	No
Enumeration area fixed effects	No	Yes	Yes	Yes	Yes	Yes
Observations	25811	25004	5248	5059	5248	5059
R-squared	0.194	0.201	0.110	0.215	0.120	0.140
N_g	2980	310	277	276	277	276
r2_w	0.194	0.201	0.110	0.215	0.120	0.140
rho	0.323	0.229	0.255	0.241	0.399	0.392

***,**,* denote statistical significance at the 1%, 5% and 10% levels. Robust t-statistics in parentheses.

Three robustness checks were carried out (not reported here): (1) restricting the sample to wave 1 and wave 2; (2) restricting the sample to maize plots and (3) discarding the bottom and top 5% of yields to deal with outliers. The results held in all subsamples: the IR was strong when yields were based on self-reporting, but disappeared when yields when crop cuts were used.

A first explanation for the finding that the existence of the IR depends on the method used to estimate yields is that both methods to estimate yields do not capture the same underlying concept. Crop cuts measure maximum attainable yields, while self-reported production is a proxy for the actual harvest. The observed inverse relation between plot size and self-reported yields is a real phenomenon if small plots are more completely harvested than large ones. This would not rule out that potential yields, as measured by crop cuts, remain constant across plots. This hypothesis is only partially upheld in the data.

Table 5 shows the results of estimating the relation between labor input per hectare and plot size. Two types of labor (family and hired labor) and agricultural activity (planting and harvesting) are distinguished. There is a strong, inverse relation between family labor (during planting and harvest) and plot size, and a positive, but less robust, relation between plot size and hired labor. This confirms that small plots are more intensively cultivated than large plots. Both findings are not novel in the literature and have been reported by Ali and Deininger (2015) in Rwanda and Larson et al. (2013) in several African countries.

Table 5: Estimation of the inverse relation between labor input per hectare and plot size

	(1) Family labor planting	(2) Hired labor planting	(3) Family labor harvest	(4) Hired labor harvest
Log of plot size (m²)	-0.382***	0.0661***	-0.311***	0.0769***
	(-40.18)	(8.30)	(-22.08)	(8.01)
Household characteristics	Yes	Yes	Yes	Yes
Plot characteristics	Yes	Yes	Yes	Yes
Labor input	Yes	Yes	Yes	Yes
Enumeration area fixed effects	Yes	Yes	Yes	Yes
Observations	24743	24743	25004	25004
R-squared	0.299	0.038	0.169	0.035

***,**,* denote statistical significance at the 1%, 5% and 10% levels. Robust t-statistics in parentheses. Full results are available upon request. All dependent variables are expressed in days/m².

While controlling for labor input during planting and harvest substantially weakens the IR, the relationship remains statistically significant (table 3, column 4). Family labor, particularly during the harvest, is correlated with yields based on self-reporting (table 3, column 4) and crop cuts (table 3, column 6), but the former correlation is much stronger than the latter. This confirms that self-reported yields are higher than crop cuts if the plot is more completely harvested. Since small plots are more completely harvested than larger ones, yields based on self-reporting are higher than crop cuts on small plots. In sum, the analyses suggest that households indeed harvest small plots more intensively than large plots, which contributes to explaining the inverse relationship between plot size and self-reported production per hectare, but does not completely explain the inverse relationship.

The second explanation for the IR is systematic measurement error in self-reported production. This can explain the IR if bias in self-reported production systematically correlates with plot size, that is, if farmers over reported production on small plots and underreported it on large plots. This can be tested directly by regressing the ratio of self-reported yields versus crop cuts on plot size (table 6), which confirms that self-reported productivity is overestimated on smaller plots. It also confirms that yields based on self-reporting increase if the plot is more completely harvested.

Based on these regressions, relative measurement error in self-reported production can be estimated as a function of plot size. This requires, however, an assumption about the parameter v, which represents the ratio of yields based on crops cuts versus unobservable 'true' yields. We show results for three values of v $(1;\ 1{,}2;\ 1{,}4)$. This choice corresponds to assuming that crop cuts are the gold standard of yield measurement (v=1) or systematically over estimate yields by 20% ($v = 1.2$) or 40% ($v = 1.4$). In order to assess the sensitivity of the results to the estimated parameters of the IR, relative measurement error in self-reported production is estimated using parameters for the core sample (left panel in figure 2), which excluded the top and bottom 5% of yields, and for the full sample (right panel in figure 2).

Table 6: Regression of the log of the ratio of self-reported yields versus crop cuts on plot size

	(1) Full sample	(2) Full sample	(3) Core sample
Log of plot size (m²)	-0.295***	-0.287***	-0.233***
	(-25.64)	(-11.99)	(-10.36)
wave	0.414***	0.319***	0.291***
	(15.58)	(5.58)	(5.48)
Labor inputs (logs)			
Family labor planting (days/ha)		0.0302*	0.0305*
		(1.69)	(1.83)
Hired labor planting (birr/ha)		0.0119	0.00698
		(0.93)	(0.60)
Exchange labor planting (days/ha)		-0.00517	-0.00490
		(-0.54)	(-0.50)
Family labor harvest (days/ha)		0.226***	0.182***
		(9.70)	(7.90)
Hired labor harvest (birr/ha)		0.0654***	0.0461***
		(4.46)	(2.80)
Exchange labor (days/ha)		0.0644***	0.0506***
		(6.38)	(5.92)
Enumeration area fixed effects	No	Yes	Yes
Household characteristics	No	Yes	Yes
Plot characteristics	No	Yes	Yes
Observations	5914	5053	4593
R-squared	0.133	0.223	0.161

***,**,* denote statistical significance at the 1%, 5% and 10% levels. Robust t-statistics in parentheses. Full results are available upon request. The core sample discard bottom and top 5% of yields based on self-reporting and yields based on crop cuts.

Independent of the choice of v, the graphs show that relative error in self-reported production decrease rapidly with plot size. Relative errors are smaller when parameter estimates are based on the core sample compared to estimates based on the full sample, but the trends are broadly similar. We will therefore only discuss the results for the core sample.

With $v = 1$, production is overreported by at least 50% on plots smaller than 80m² and underreported by more than 25% on plots larger than 1500m². For larger values of v, overreporting on small plots increases, while underreporting on larger plots decreases. The critical threshold of plot size below which over reporting of production changes to underreporting of production depends on the choice of v. This critical threshold equals 445m², 992m² and 2000m² for $v = 1, 1.2$ and 1.4, respectively.

Figure 2: Relative measurement error in self-reported production in function of plot size (m²)

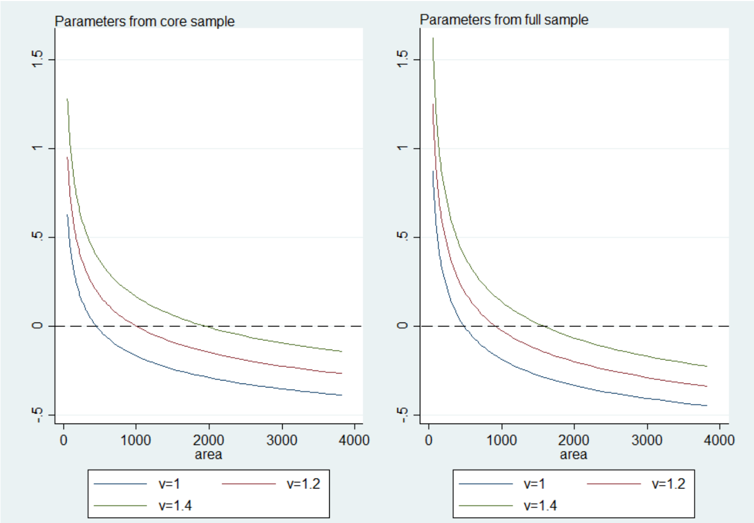

With an additional assumptions about 'true' yields, absolute measurement errors in self-reported production can be estimated. Assume that average 'true' yields are 1000 kg/ha. Under this assumption, 'true' production on a plot of 100m² is 10kg and over reporting by 50% implies an overestimation by 5kg. With the same assumptions, 'true' production on a plot of 2500m² is 250kg, but is underreported by 80kg. While an overestimation by 5kg on small plots seems reasonable, the large underestimation on larger plots appears less credible. It is worth pointing out, however, that a similar pattern has been observed for self-reported land area: farmers tend to overestimate the area of small plots, but underestimate the area of larger plots (Carletto, Gourlay, et al., 2015). Moreover, the order of magnitude of over/underestimation is similar to our findings. For instance, Carletto, Gourlay, et al. (2015) report that farmers overestimate plot size by 103% relative to GPS measurement on plots smaller than 0.5 acres (2000 m²) and underestimate it by 33% on plots larger than 5 acres (20,000 m²).

5. Conclusion

This paper examined the decades old puzzle of the inverse relation between plot size and land productivity in developing countries. It argues that the relationship is an artifact of systematic measurement error in self-reported production. This explanation has not yet been explored in the literature. Yet, most studies on the inverse relationship rely on self-reported production to estimate yields. We use an alternative measure: crop cuts. Crop cuts are an objective approach to measure yields, based on randomly sampling a subplot within a plot and cutting and weighing the harvest within the subplot. Crop cuts are sometimes considered the gold standard of yield measurement and are frequently used in specialized agricultural surveys.

Drawing on a dataset from Ethiopia, which implemented both methods of yield measurement on the same plots, we show that the inverse relationship is strong when yields are based on self-reported production, but disappears when crop cuts are used. We even observe a positive correlation between plot size and crop cuts.

In theory, these results could be explained by differences in labor input between large and small plots since crop cuts measure the potential harvest (i.e. just before the harvest) and self-reported production the actual harvest (i.e. just after the harvest). If small plots are more completely harvested than large plots, an inverse relation would only be observed for self-reported yields and not for crop cuts. Although we did find evidence of an inverse relation between plot size and family labor, the inverse relationship remained significant after controlling for labor input.

This leaves systematic measurement error in self-reported production as the remaining explanation for our findings. Nobody will dispute that self-reported production in household surveys is always measured with substantial error for various reasons, including the long recall period. Yet, this error is invariably assumed to be random. Random errors never induce spurious relations. The IR can only be explained by measurement error if the errors are systematically correlated with plot size. Our findings suggest that households over report production on small plots and underreport it on larger ones. Moreover, the systematic bias is substantial. Our best estimates indicate that production needs to be over reported by 40% to 100% on plots of 100m² and underreported by 10% to 35% on plots of 3000m² to explain the IR.

This paper does not reveal why farmers systematically over and underreport production. One may speculate that these errors are related to heaping or rounding errors, which occur frequently in survey data. A small, upwards rounding error in production on a small plot translates into a substantial error in yields when production is divided by the small plot size. On larger plots, small rounding errors are a lesser concern because a small rounding error in production will not be amplified to a large error in yields when production is divided by a large plot size. This asymmetry between small and large plots causes decreasing measurement error in yields with increasing plot size, which could generate an inverse relation. However, it does not yet explain why farmers systematically round upwards on small plots. This is an interesting avenue for future research.

Our results may be partially driven by the small plots size in Ethiopia, where the average plot size was 0.13ha. Because plot sizes are so small, small measurement errors in self-reported production translate into large measurement errors in yields. These errors may be less important when average plot size is larger or when total production is aggregated at farm level, which is often done when the inverse relation between farm size and productivity is studied. Consequently, systematic measurement error in production is especially troublesome when estimating the inverse productivity plot-size relationship, but its effects may be less pronounced for the inverse productivity farm-size relationship.

Although our analyses should be replicated in other settings, the rejection of a negative correlation between plot size and crop cuts gives nevertheless some solace to the agricultural economists that find the far-reaching policy implications of the IR disturbing. The hypothesis of missing markets as the main explanation for the inverse productivity farm size relationship has been discredited because it fails to explain productivity differences across plots within the same household (Assunção & Braido, 2007). Attributing the inverse productivity plot-size relationship to measurement error thus reinvigorates the more conventional explanation of missing markets for the inverse productivity farm-size relationship. From a policy perspective, this implies that reducing frictions in land and labor markets will increase agricultural output. The rejection of the inverse productivity plot-size relationship has also important implications for land consolidation programs since the IR has frequently been cited as an argument against land consolidation. Although the positive correlation between plot size and crop cuts is weak in our data, we can at least claim that land consolidation will not necessarily reduce agricultural production.

References

Aguiar, M., & Bils, M. (2015). Has Consumption Inequality Mirrored Income Inequality? American Economic Review, 105(9), 2725-2756. doi:doi: 10.1257/aer.20120599

Ali, D. A., & Deininger, K. (2014). Is there a farm-size productivity relationship in African agriculture? evidence from Rwanda.

Ali, D. A., & Deininger, K. (2015). Is There a Farm Size–Productivity Relationship in African Agriculture? Evidence from Rwanda. Land Economics, 91(2), 317-343.

Assunção, J. J., & Braido, L. H. B. (2007). Testing Household-Specific Explanations for the Inverse Productivity Relationship. American Journal of Agricultural Economics, 89(4), 980-990. doi:10.1111/j.1467-8276.2007.01032.x

Barrett, C. B. (1996). On price risk and the inverse farm size-productivity relationship. Journal of Development Economics, 51(2), 193-215. doi:http://dx.doi.org/10.1016/S0304-3878(96)00412-9

Barrett, C. B., Bellemare, M. F., & Hou, J. Y. (2010). Reconsidering conventional explanations of the inverse productivity–size relationship. World Development, 38(1), 88-97.

Beegle, K., De Weerdt, J., Friedman, J., & Gibson, J. (2012). Methods of household consumption measurement through surveys: Experimental results from Tanzania. Journal of Development Economics, 98(1), 3-18.

Benjamin, D. (1995). Can unobserved land quality explain the inverse productivity relationship? Journal of Development Economics, 46(1), 51-84. doi:http://dx.doi.org/10.1016/0304-3878(94)00048-H

Bhalla, S. S., & Roy, P. (1988). Mis-Specification in Farm Productivity Analysis: The Role of Land Quality. Oxford Economic Papers, 40(1), 55-73. doi:10.2307/2663254

Blarel, B., Hazell, P., Place, F., & Quiggin, J. (1992). The economics of farm fragmentation: evidence from Ghana and Rwanda. The World Bank Economic Review, 6(2), 233-254.

Boumans, M. (2015). Science Outside the Laboratory: Measurement in Field Science and Economics: Oxford University Press.

Carletto, C., Gourlay, S., & Winters, P. (2015). From guesstimates to GPStimates: land area measurement and implications for agricultural analysis. Journal of African Economies, ejv011.

Carletto, C., Jolliffe, D., & Banerjee, R. (2015). From tragedy to renaissance: improving agricultural data for better policies. The Journal of Development Studies, 51(2), 133-148.

Carletto, C., Savastano, S., & Zezza, A. (2013). Fact or artifact: The impact of measurement errors on the farm size–productivity relationship. Journal of Development Economics, 103(0), 254-261. doi:http://dx.doi.org/10.1016/j.jdeveco.2013.03.004

Carroll, R. J., Ruppert, D., Stefanski, L. A., & Crainiceanu, C. M. (2012). Measurement error in nonlinear models: a modern perspective: CRC press.

Carter, M. R., & Wiebe, K. D. (1990). Access to capital and its impact on agrarian structure and productivity in Kenya. American Journal of Agricultural Economics, 72(5), 1146-1150.

Chayanov, A. V. (1926). AV Chayanov on the theory of peasant economy: Manchester University Press.
Collier, P., & Dercon, S. (2013). African Agriculture in 50Years: Smallholders in a Rapidly Changing World? World Development.
De Groote, H., & Traoré, O. (2005). The cost of accuracy in crop area estimation. Agricultural Systems, 84(1), 21-38. doi:http://dx.doi.org/10.1016/j.agsy.2004.06.008
Eswaran, M., & Kotwal, A. (1986). Access to capital and agrarian production organisation. The Economic Journal, 96(382), 482-498.
Feder, G. (1985). The relation between farm size and farm productivity: The role of family labor, supervision and credit constraints. Journal of Development Economics, 18(2), 297-313.
Fermont, A., & Benson, T. (2011). Estimating Yield of Food Crops Grown by Smallholder Farmers. International Food Policy Research Institute. June.
Gibson, J., & Kim, B. (2007). Measurement Error in Recall Surveys and the Relationship between Household Size and Food Demand. American Journal of Agricultural Economics, 89(2), 473-489. doi:10.2307/4492824
Hausman, J. (2001). Mismeasured variables in econometric analysis: problems from the right and problems from the left. The Journal of Economic Perspectives, 15(4), 57-67.
Jerven, M. (2013). Poor numbers: how we are misled by African development statistics and what to do about it: Cornell University Press.
Lamb, R. L. (2003). Inverse productivity: land quality, labor markets, and measurement error. Journal of Development Economics, 71(1), 71-95. doi:http://dx.doi.org/10.1016/S0304-3878(02)00134-7
Larson, D. F., Otsuka, K., Matsumoto, T., & Kilic, T. (2013). Should African rural development strategies depend on smallholder farms? An exploration of the inverse-productivity hypothesis. Agricultural Economics, n/a-n/a. doi:10.1111/agec.12070
Pašakarnis, G., & Maliene, V. (2010). Towards sustainable rural development in Central and Eastern Europe: Applying land consolidation. Land Use Policy, 27(2), 545-549.
Sen, A. K. (1962). An aspect of Indian agriculture. Economic Weekly, 14(4-6), 243-246.
Taffesse, A. S., Dorosh, P., & Asrat, S. (2011). Crop production in Ethiopia: regional patterns and trends.