

Doss, Cheryl R.

Working Paper

Women and agricultural productivity: What does the evidence tell us?

Center Discussion Paper, No. 1051

Provided in Cooperation with:

Economic Growth Center (EGC), Yale University

Suggested Citation: Doss, Cheryl R. (2015) : Women and agricultural productivity: What does the evidence tell us?, Center Discussion Paper, No. 1051, Yale University, Economic Growth Center, New Haven, CT

This Version is available at:

<http://hdl.handle.net/10419/147535>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

**ECONOMIC GROWTH CENTER
YALE UNIVERSITY**

P.O. Box 208269
New Haven, CT 06520-8269

<http://www.econ.yale.edu/~egcenter>



Economic Growth Center Discussion Paper No. 1051

**Women and Agricultural Productivity: What Does the Evidence Tell
Us?**

Cheryl Doss
Yale University

Notes: Center discussion papers are preliminary materials circulated to stimulate discussion and critical comments.

This paper can be downloaded without charge from the Social Science Research Network Electronic Paper Collection: <http://ssrn.com/abstract=2682663>

Women and Agricultural Productivity: What Does the Evidence Tell Us?

Cheryl Doss

Abstract: Should agricultural development programs target women in order to increase productivity?

This paper reviews the extensive literature on men's and women's relative productivity in agriculture, most of which concludes that controlling for access to inputs, plot and farmer characteristics, there are little or no gender gaps in productivity. In addition, the paper identifies the many challenges to disentangling individual level productivity. Most of the literature compares productivity on plots managed by women with those managed by men, ignoring the majority of agricultural households in which men and women are both involved in management and production. The empirical studies which have been done provide scant evidence for where the returns to project may be highest, in terms of who to target. Yet, programs that do not consider the gendered responsibilities, resources and constraints, are unlikely to succeed, either in terms of increasing productivity or benefitting men and women smallholder farmers.

Keywords: agricultural productivity, gender, smallholder farming, developing countries, agricultural policy

JEL Codes: 13, O33, O20, Q12

This paper is based on research funded in part by the Bill & Melinda Gates Foundation. The findings and conclusions contained within are those of the authors and do not necessarily reflect positions or policies of the Bill & Melinda Gates Foundation. Diksha Arora provided excellent research assistance. The paper benefitted from comments from Douglas Gollin, Deborah Rubin, and Christopher Udry.

1. Introduction:

To what extent should agricultural development investments explicitly target women? Is it possible to use investments in agricultural technology to redress gender imbalances and to reduce poverty? Would a women-focused agricultural investment strategy yield a double dividend – increasing gender equity at the same time as increasing overall productivity? These and other questions have become pressing issues for development policy and strategies. This paper will discuss the available evidence on these questions and will also discuss both the measurement and conceptual challenges that have confronted efforts to answer these important policy questions.

The development literature abounds with claims about the benefits of targeting agricultural investments to women, especially in sub-Saharan Africa. These claims take many forms, but in general it is argued that increasing women's agricultural productivity is the key to increasing overall agricultural productivity, empowering women, and reducing poverty.

In general, the arguments for targeting women can be grouped into two main strands. One strand focuses on the productive potential of women farmers. A generic argument here would be that women are heavily involved in agricultural production in the developing world – and especially in Africa – and that they have been left out of many development efforts. In consequence, this might imply that there are very high returns to targeting current investments to women – with these returns showing up in the form of increased aggregate agricultural production as well as in the income earned by women.

A second (and not mutually exclusive) strand of argument argues that women represent an important class of beneficiaries of agricultural development efforts – and that their needs have frequently been neglected by programs that focus on productivity increases. Because many poor women are farmers, and many poor farmers are women, there are reasons to direct agricultural development towards this group. The importance of women as beneficiaries is increased by the instrumental roles of women with respect to child health, nutrition, and education. Improving the well-being of women and offering them expanded opportunities will not only increase their own welfare but also has the potential to create positive effects on the next generation.

Both of these strands of argument incorporate a range of hypotheses, and both rest on plausible and internally consistent theories of change. In essence, they are arguments that the social rates of return on agricultural development investments are higher when those investments are targeted to women.

The empirical basis for these arguments is not nearly as well developed, however. No consensus has emerged on the social benefits of targeting research or other agricultural interventions to women. The claim of efficacy relies on several steps in a complex causal chain: first, that investments can selectively drive up women's productivity; second, that these increases in women's productivity will actually succeed in producing benefits for women (and perhaps also children); and third, that the resulting social rates of return for these investments are higher than for other development investments.

Many voices in the development community take it as a given that these causal links are satisfied, and donors have increasingly required that gender issues must be addressed in projects and proposals. Other voices continue to express skepticism about a women-focused strategy in agriculture – or at least suggest that there may be tradeoffs associated with targeting interventions to women. The empirical analyses of these issues has been rather sparse.

A recent set of papers, using sex-disaggregated data that has recently become available, has examined some of the oft-cited statistics about women's role in agriculture. Doss et al. (2015) examine the claim that women own 1-2% of the world's property and show that for Africa, although women own considerably less land than men, they own more than 1-2% of the land that is owned by individuals (as opposed to state or corporate land). Kieran et al (2015) reach similar conclusions for Asia. Palacios-Lopez et al. (2015) find that, in the six African countries for which there is sufficient data, women provide 40% of the labor for crop agriculture, a lower figure than the number often cited (which is an unrealistic 60-80%). They do find that the patterns vary widely both across and within countries. Finally, Doss (2014) analyzes the claim that women produce 60-80% of the world's food. She argues that there is no evidence to support this claim, and that, given women's lower access to land and other inputs, and their responsibilities for household work, it would be surprising if they produced most of the food.

A related literature, with a much longer history, focuses on the relative productivity of men and women in agriculture. While most studies find that women produce less on their plots of land, much of the effort has focused on demonstrating that this is because women have lower quality land, less access to fertilizer and other inputs, and receive less credit and extension support. This literature, discussed in detail below, suggests that women have the potential to have higher productivity, given the necessary resources.

After reviewing much of this literature, The FAO State of Food and Agriculture Report in 2010-11 claims, "If women had the same access to productive resources as men, they could increase yields on their farms by 20-30 percent. This could raise total agricultural output in developing countries by 2.5-4 percent, which could in turn reduce the number of hungry people in the world by 12-17 percent." These are plausible estimates, but in this era where the gold standard of evidence is randomized control trials, it is important to remember that these numbers are calculated using the estimated production functions for women assuming that they used the same levels of inputs as the men. They are not the result of programs providing women with the same level of inputs as men to see how their production would increase.

Beyond the statistics, some of the discussion around women-focused agricultural development involves rather blurry conceptual thinking, both from advocates and skeptics. As this paper will argue below, it is not clear what is meant by "women's productivity." This is a difficult concept to define, let alone to measure. In the same vein, what would it mean to identify food as being "produced by men" or "produced by women." Are these useful ways to think about the gendered structure of agriculture and food systems? Across the world, most farms are operated by families that include both men and women, so how can we disentangle their individual contributions?

This paper will argue that much of the academic literature on women in agriculture has gone in directions that are not helpful in shedding light on the central policy questions. Similarly, the policy debate has often highlighted the wrong issues. The literature offers little insight into the key question – which is, simply put, where the returns to development investments are highest. The paper will also argue that however investment priorities are set, gender is an essential analytic category for thinking about the impacts of agricultural development investments. This is true regardless of whether or not women are targeted as the users of technology or even as beneficiaries. Farming and food preparation are deeply gendered activities. The impacts of agricultural technologies and development interventions

are necessarily filtered through the gendered patterns of agricultural labor, household enterprises, family food consumption decisions, and social structures. Agricultural technologies are not, in general, gender neutral. Thus, a gender lens is essential for assessing the effectiveness and impact (whether *ex ante* or *ex post*) of an agricultural technology or intervention.

From a policy perspective, many development institutions and research organizations have adopted explicit emphasis on targeting agricultural development to women. In many cases, this is part of a growing realization that women are disadvantaged in many dimensions within developing countries. This paper will not summarize the history of efforts to target women, nor will it repeat the oft-cited justifications. It is worth noting, however, that in spite of these stated goals, most interventions still do not have an explicit gender orientation; and in many cases, interventions do *not* specify a gender target. This implies a default in which many or even most interventions actually target men, with a smaller and separate set of interventions targeted at women farmers.

So what does the evidence show? Is there any evidence that would support or refute the basic claim about social returns to investments that target women, whether as producers or as beneficiaries? Unfortunately, the literature has very little to offer on these big questions. In part, this reflects the difficulty of constructing appropriate and rigorous tests of these hypotheses. There is no experimental framework in which this question can easily be asked at the level of an institution's development portfolio. At the project level, there is abundant information about whether specific interventions create benefits for women. But this project-level information does not necessarily address the larger strategic questions about how research institutions and donors should shape their investment portfolios.

This paper will review the existing literature on women's productivity in agriculture. Section 2 starts by discussing how we measure agricultural productivity. Building on these definitions, section 3 asks how we can separate women's agricultural productivity from men's and identifies some of the conceptual challenges facing this literature. Section 4 reviews the empirical literature, considering each of the approaches that have been used and the evidence provided. Section 5 briefly outlines the challenges of considering women's agricultural productivity in systems where men and women farm jointly. Finally, section 6 concludes.

2. Measuring Agricultural Productivity

The first challenge in comparing men's and women's productivity in agriculture lies in the measurement of agricultural productivity. A related conceptual issue is how to measure the separate productivity of men and women, especially when they are typically working alongside one another as part of complex family enterprises. As a practical matter, it is difficult to measure productivity at the level of a plot or a farm – and it becomes even harder when we try to disentangle the productivity of individuals within the farm household, especially when gender roles and norms assign women and men to different tasks.

Agricultural productivity is normally measured in a several different ways. Each approach has advantages and disadvantages and different data requirements. These approaches do not readily lend themselves to measuring separately the productivity of men and women, so additional assumptions are needed to arrive at sex-disaggregated measures of productivity. The additional assumptions (e.g., on functional specification of the agricultural production technology) in turn have the potential to affect the results and conclusions concerning women's and men's productivity levels.

The simplest way to think about productivity is that reflects a relationship between inputs and outputs. The two most widely used measures are so-called partial productivity measures – output per unit of a single input. Land productivity is defined as output per unit of land; labor productivity is output per unit of labor effort. In both cases, output may be defined in either physical quantities (such as kilograms of grain) or economic value. Labor input is normally defined in terms of hours of time, but in some cases is measured in value terms – for example, when hired labor is aggregated with family labor.

Land Productivity

The most common measure of agricultural productivity is a simple measure of crop yield, which is the physical quantity of output per unit of land, for a single crop. This approach is comparable across areas. It works best when farmers plant one crop on a plot of land, making it conceptually straightforward to calculate the quantity of output on a randomly selected sub-plot. The measurement itself is not trivial; accepted practice is to follow complex protocols for conducting “crop cuts” at harvest time, with concomitant protocols governing technical issues such as the moisture content of the crop at the time when the output is weighed. In some instances, however, yield is measured through farmer recall or other estimates of how much output was produced, along with measures of the size of the plot.

Yield measurements can be complex in areas where there is multiple cropping per year on a single plot of land or where several crops are simultaneously planted. In the case of multiple cropping in sequence, a higher annual output may sometimes be attained from two different harvests per year – even though the output in each season is lower than for single cropping. The picture becomes even more complicated when different crops are grown in different seasons, or when multiple crops are cultivated simultaneously on the same field, as is common in many African production systems.

When multiple crops are grown on a single plot or when we are interested in farm-level productivity rather than a plot with one crop, there are three possible approaches. The simplest, used often with yield, is to allocate the input of land or labor to each of the types of output. For example, if an acre plot is intercropped equally with maize and beans, then half an acre would be considered to be planted with maize and half an acre with beans and the yields calculated accordingly. A second option is to treat intercropped plots as though it were a separate crop for the purposes of analysis. Thus, we might analyze plots planted maize, those planted to beans, and those intercropped with maize and beans.

The third option is to aggregate across crops. This is typically done by considering the value of output, usually in farm-gate prices. This approach is conceptually very clear and has the advantage of allowing comparisons across crops and cropping systems. However, the value measures are quite noisy and may vary widely across time and location. The same quantity of physical output may receive different prices in different years (or even on different days), depending on total supply and the availability of local markets. Prices vary widely across farms and farmers, due in part to differences in location.

Most large-scale surveys typically find large differences in yield across plots of land, even within geographically narrow areas. These may be due to a range of factors. There may be differences in plot-level characteristics, such as the soil quality or availability of irrigation. Different plots may also experience different shocks within a growing season, such as weather or exposure to pests or disease. Plots may be differentially located with respect to infrastructure, including roads and markets. Other differences may result from the choices of farmers, such as which variety of seed to use, whether to use fertilizer, and particular management practices. Finally, the characteristics of the farmer may be important, such as farming experience or level of education. The sex of the farmer may be correlated with all, some, or none of these factors.

This means that comparisons of yield for men and women farmers are difficult both to make and to interpret. Where men and women own and/or manage separate plots, the yield differences could in principle be due to the managerial abilities of the men and women farmers, or to their physical strength or other attributes. But they may also be due to differences in the plots or other to some of the many complicated choices that households make about how to allocate resources and prioritize investments across a range of agricultural and non-agricultural activities. Thus, a comparison of yields on male- and female-owned plots, or yields attained by male- and female-headed households, has little or nothing to tell us about the potential returns to investments in agriculture that target women.

Labor Productivity

Labor productivity is similar to yield, but it measures the output per unit of labor input. Again, when comparing labor productivity across individual farms, a number of factors will play a determining role. The quality and location of land will be important; and the type and quality of labor will also affect output. A day of labor will have a different marginal contribution to the final level of output, depending on the characteristics of the person providing the work – varying by age, sex, health status, etc. Some workers will put in longer days and more effort. Simply measuring labor input is challenging, even before disaggregating by sex. In addition, the returns to a day of labor will vary across the agricultural season and activity; the contribution of a day spent weeding may be different from the contribution of a day spent clearing the field and it is difficult to disentangle these effects.

When calculating labor productivity, again output measures are often calculated in terms of value, so that they can incorporate a range of crops. As before, it is important to note that in any given growing season, agricultural production is highly dependent on local weather shocks, pest problems, and so forth. This generates a lot of statistical “noise” in production data, which means that labor productivity is not calculated with much accuracy in any season. To distinguish differences in labor productivity between two farms or two individuals, we would typically expect to need lots of observations to extract “signal” from the noise – even if everything else were held exactly equal.

Numerous studies have tried to estimate labor productivity for men and women separately, and new household survey data such as the LSMS-ISA studies provide fairly detailed estimates of labor use on each plot by different individuals. But it is difficult to extract meaningful measures of individual productivity from these data, and even then it is hard to know how to interpret measures of labor time. Much agricultural labor is carried out jointly with other household activities. How should we count hours that a woman spends working on the farm while supervising or caring for her children? Or the time that she spends tending backyard livestock while carrying out other activities around the homestead? Indeed, how should we account for cultural views of labor that may not view certain of a woman’s tasks as “work”? Studies of labor productivity typically blur over these problems, and comparisons of men’s labor productivity with women’s labor productivity are consequently problematic.

Profitability

Economists often prefer to use measures of profitability rather than the simple measures of yield or labor productivity. Profitability allows for the fact that both land and labor (and indeed other inputs) may vary. A farmer may produce at a higher level of yield by using more inputs – but this may or may not be profitable. Incorporating costs, such as the cost of inputs, provides a more robust measure of overall

productivity. Measures of profitability use monetary values to sum across crops and net out the cost of inputs.

Estimating profitability requires information not only on the costs of outputs, but also on the costs of inputs. Value added measures are calculated by subtracting the cost of purchased inputs, such as fertilizer and seeds, from the value of output.¹ Prices may vary a great deal from one farm to the next, due to “last mile” transport costs and quality differences. Different farmers may purchase their inputs (and sell their outputs) at different moments during the season, with corresponding variation in prices. The costs of factors of production (typically, land, labor and capital) are not included in these measures of agricultural profitability. Identifying the costs of these latter factors of production can be quite challenging, especially when the markets for them are thin. In addition, the opportunity cost of labor varies by individual and across the season.

From an economic perspective, efficiency does not actually imply that farmers should be maximizing production and certainly not that they should be maximizing the production of a particular crop. Instead, economists understand efficiency as closely related to profit maximization. Theory suggests that farmers will seek to maximize the profitability of their farms, rather than simply maximizing the total level of output of one or more crops. The farms may include both crops and livestock as well as other income generating activities. The distinction implies that measures such as yield and labor productivity can give a misleading impression. For example, under many plausible circumstances, increasing yield may actually decrease profits; for example, when fertilizer and improved seeds are relatively expensive and the farm gate price of outputs is low. This could be a result of remote location and limited access to markets.

For this reason, economists tend to like measures of productivity that account for many different inputs and factors of production. Measures of productivity that account for multiple factors of production are typically called “multi-factor productivity” or “total factor productivity” (TFP) measures. As generally calculated, agricultural measures of TFP normally account simultaneously for the use of land, labor, and capital. TFP measures can be understood as comparing aggregated output to aggregate inputs. Calculating TFP raises challenges of its own and requires a number of assumptions (e.g., about the ways in which inputs can be aggregated). Unlike yield differences across plots, a measure of total factor productivity would at least address the problem that some of those differences emanate from differential input use and differences in quantities of land, labor and capital.

But as with measures of yield and labor productivity, TFP exhibits wide heterogeneity across plots, due to factors such as localized weather. Calculations are also quite sensitive to assumptions (such as the form of the production function). Often TFP is used to analyze changes over time. We say that TFP has increased when the index of outputs increases faster than the index of inputs. Yet, for differences in TFP to be convincing, we need many data points, so that we can be sure that the variations in output due to weather and other short term factors are not the driving factors. This might require multiple observations over time, or lots of observations of the same farmers under different circumstances. Comparing profitability or TFP for men and women farmers is thus very demanding in terms of data. These comparisons assume that inputs and outputs are measured well (and with appropriate allowances for quality differences; e.g., for land and machines). It is also essential that inputs and outputs are measured

¹ To calculate the profitability of particular crops requires this information at the crop level; the fertilizer used on the particular crop. Again, this is challenging when intercropping is practiced; e.g. how should the cost of the input be allocated across crops?

equivalently for men and women, so that (for instance) women's labor time is not overstated when they are engaged in other activities concurrent with farming.

3. How do we distinguish men's and women's productivity?

Given all the difficulties in measuring productivity, how can we distinguish the productivity of men and women in agriculture? This is a topic that has received quite a lot of attention – though as noted above it is not clear how and whether this relates to the rates of return on investments targeted to women.

In comparing productivity of men and women, a key challenge is to assign household farm production to individuals. How do we separately measure men's and women's agricultural productivity?² This section first discusses conceptual and measurement issues and then turns to specific methodological approaches.

Conceptual issues and measurement

For comparing men's and women's agricultural productivity, one approach that has been widely used is to consider the "household farm enterprise" as the production unit and to compare the productivity of different households, making a distinction between male- and female-headed households. While this approach is relatively simple, it ignores the contributions that women make to farms in households headed by men (and conversely the contributions that men make to farms in households headed by women). This not only ignores the joint structure of most households that are headed by couples; but it also introduces a possibly spurious comparison between households that differ in many other ways. Cultural norms frequently dictate that a household will not be defined as female-headed if there is an adult male present, regardless of the number of women. This means that female-headed households almost invariably have fewer adult members (and no male members), whereas male-headed households are generally larger and have both men and women. The two household types are clearly not comparable. A comparison of female-headed households and male-headed households does not tell us anything about men and women farmers; instead, it is a comparison of two different household structures. Thus, the literature comparing male- and female-headed households will not be included in the empirical review in this paper.³

A second approach is to consider the sex of the farmer of a particular plot or set of plots. This approach may allow for multiple farmers in a given household, and it has become feasible with more recent household surveys that ask questions about the farmer for each plot controlled by a household. Typically, output from a plot is assigned to an individual farmer, defined either as the person who "manages" the plot or the person who "controls output" from the plot. The output of plots managed by women is then treated as women's production. This approach has largely replaced a second option – which was often the only possible approach for older data sets – namely, to assign the output of particular crops to different members of the household. For example, one might assume that cultural norms in a particular context mean that only men grow cocoa, or only women grow millet. The household output of these crops would then be assigned to men and women based on these

² For a much more detailed discussion of this issue, specifically on how to identify how much food is produced by women, see Doss (2014).

³ In spite of the problems of this comparison, it has been widely used in the literature, and researchers continue to characterize this as a form of gender analysis. The CGIAR's new minimum standards for sex-disaggregated data collection explicitly reject this.

characterizations of gendered patterns of farm production. This approach is particularly problematic if we are trying to follow changes over time, since the patterns of who grows what is likely to change.

This discussion begs the question of how a farmer is defined. In some cultural contexts, women may be seen as farm helpers, but not farmers; or they may be viewed as only growing food in the kitchen garden or homestead (Twyman et al, 2015). In other instances, women are identified as independent farmers, either by plot or crop. The person identified as the farmer may be the owner of the plot, the manager of the plot, or the person providing the day to day labor for the plot or crop. Because women are less likely to own the land or to have independent rights to the land, they may not be considered the farmer. If we misidentify the farmer, then our measures of men's and women's productivity will not be useful.

It is important to note that plot-based and crop-based measures of output almost always omit household activities related to animal agriculture. This is potentially a very important omission: by most estimates, livestock-related activities account for a large fraction of total agricultural production in almost all farming systems – and often for even larger fractions of cash income. For example, FAO (2009) reports that livestock contribute 40 percent of the global value of agriculture output. Livestock also play important roles as household assets, including as stores of value and as a form of insurance. It is a non-trivial omission to exclude animal agriculture from calculations of men's and women's agricultural productivity. Significant amounts of household labor and other inputs are often devoted to animal husbandry, and significant fractions of household earnings are derived from the output of animal agriculture. Moreover, animal care is often highly gendered – with women responsible for dairying and smallstock in many societies, while men are typically responsible for ranching and herding. This implies that in different farming systems, the omission of animal agriculture may have important implications for the measurement of men's and women's relative agricultural productivity.

Methodological approaches

The previous section argues that there are conceptual problems with assigning output to individual men and women. But as a purely statistical matter, when a particular plot or set of activities can be assigned to a particular man or woman, along with the associated output, it is relatively easy to attribute the productivity to that person. In a regression analysis – whether of yield, profit, or efficiency – we simply categorize the farmer as male or female, using a binary dummy variable. If the coefficient on this dummy variable is statistically significantly different from zero, then we would claim that there are gender differences in the corresponding left-hand variable. A crude analysis would regress the outcome variable only on the sex dummy; a more sophisticated analysis would include a number of control variables on the right-hand side, perhaps including input levels and a vector of other household characteristics. If the sex dummy remains significant in this regression, the argument would be that the gender of the farmer has an impact on productivity above and beyond any differences in inputs or characteristics. A similar option is to estimate separate equations for men's productivity and women's productivity and then test whether they are significantly different.

Recent approaches have gone a step further, using decomposition techniques, such as the Oaxaca-Blinder decomposition (discussed in more detail below). These make it possible to assess how much of the differences in productivity between men and women can be attributed to differences in use of inputs and how much is unexplained and therefore attributed to gender.

These approaches cannot consider the productivity of different individuals on a farm that is jointly managed or to consider the contributions of women on farms that are managed by men. We return to this issue below; but first, it is useful to consider some of the underlying challenges with these estimations.

One issue relates to crop choice. Productivity measures usually take the choice of crops as given, and ask, for example, whether men and women are equally productive as maize farmers. If in each household, the decision of which individuals cultivate which crop on which plot were to be made randomly, then the productivity comparisons would require nothing more than a straightforward application of statistics. In reality, however, this is not the case. Plots are not assigned randomly to men and women; nor are crops assigned randomly to plots. If households systematically assign the best plots to men, and if these plots are used to grow maize primarily for the market; whereas the worst plots are assigned to women to grow millet intercropped with beans for home consumption, how do we make comparisons? What would it mean to claim that men are more productive at growing maize than women are at growing millet or beans? Men and women farmers face different access to inputs, information and markets, and thus they will make different choices about what to grow. These choices may be made jointly, as part of a household strategy, or they may be made separately and with little discussion or cooperation. Decisions of what to grow (and how to grow it) may reflect profit-maximizing strategies, or they may be driven by social norms and cultural dictates.

The question of crop choice also is related to the extent to which men and women are producing for the market or home consumption. When producing partly or entirely for home consumption, farmers may not be maximizing profits (or production) but may instead be pursuing a range of more complex objectives. When there is limited access to markets, with a high wedge between prices farmers receive as producers and the amount that they pay as consumers, it can be economically rational for them to produce lower-value crops for home consumption rather than more valuable crops for the market.

Gender conditions many other aspects of crop choice and cultivation practices. Men and women may have different access to markets to sell their produce. They may value different traits for the same crop (e.g., taste and cooking characteristics for home-consumed crops, versus yield or cost advantage for marketed crops). Men and women may have different access to tools and machines, and this may affect both their choices of crops and the techniques of cultivation.

Another issue that arises in estimating the coefficient on the sex dummy variable is that input use is also endogenous. To pursue further the example above, men and women growing maize on different plots will rationally use different types and quantities of inputs. Fertilizer and other agricultural chemicals are not randomly assigned to men's plots and women's plots; these and other inputs may vary in systematic ways with the sex of the farmer. A number of factors will affect input use, including prices, transport costs (e.g., to get the inputs to the farm), or the units in which they are available for purchase. And again, gender conditions all of these factors. Men and women may also have different access to information about technologies and markets, both through extension services and through their own informal networks.

While showing that there are productivity differences between husbands and wives does not necessarily mean that the household as a whole is not optimizing, we would expect that the marginal product of all of the inputs, including land and labor, would be the same across plots. If this is not the case, then it suggests that there are some missing markets within households. A variety of reasons may limit the allocation of resources between spouses, including information failures, power dynamics or social norms.

There are a range of other ways in which gender may influence the choices that farmers make about what to grow and how to grow it. In general, women have smaller – and often poorer quality – plots of land. Thus, the scale of farming may differ between men and women. Access to markets may differ by gender; some markets are not understood as appropriate places for women, and the access to transportation may vary. (For example, in some places, social norms prohibit women from riding bicycles.) Some evidence suggests that there may be systematic gender differences in the prices received by men and women for the same output. Women may sell more of their produce at the farm gate, if their time is limited by household responsibilities that reduce their ability to travel to the market. Women may also sell at the farm gate because they are selling smaller quantities and it is not worthwhile to make the trip to the market. There may be social strictures that make it more difficult for women to bargain effectively with male traders for higher prices. Certainly, not all of these factors are present in every community; there are women who ride bicycles, bargain well in the market, and travel long distances to sell their produce. But it is important to consider the ways in which gender may mediate many of these market activities and thus choices regarding agricultural production.

4. Empirical Estimates of Gender Differences in Agricultural Productivity

In spite of all of these challenges, there is a large and growing empirical literature that claims to estimate the gender differences in agricultural productivity. One motivation for much of this literature is to ask whether women have the potential to be as productive as men, given the same access to productive resources. The question is whether women are inherently less productive farmers (or perhaps more productive farmers), or whether they simply face a different range of constraints than those facing men.

A range of empirical approaches have been used to estimate the gender differences in agricultural productivity. This literature on gender differentials in agricultural productivity can be broadly categorized as using five different approaches: estimating production and profit functions, calculating the marginal productivity of labor from production functions, computing technical inefficiency using stochastic production frontier methods, applying Oaxaca-Blinder decomposition to calculate endowment and structural effects on productivity differentials, and, finally, employing technology adoption as a proxy. To some extent, these different approaches are asking somewhat different questions and thus, it is useful to consider this literature by the approach that is used. In this section, I discuss each of the approaches and the findings that come out of these studies.

Estimating production and profit functions

In this approach, regression analysis is used to estimate the contribution of various inputs to the attained output. This lets us move beyond a simple comparison of men's and women's average output to ask about men's and women's levels of output conditional on inputs. Output can be measured either as yield or as value.

The intuition behind these regressions is that they establish whether there is a significant relationship between the sex dummy variable and yield or profit even after controlling for observable input differences and other plot or farmer characteristics. Although there are numerous studies and many conflicting results, it is fair to say that in general this literature finds little or no significant difference in men's and women's agricultural productivity after controlling for access to inputs, plot, farmer, and

household characteristics. This is usually interpreted to mean that women would be equally productive as men if they were given the same access to land, education, information, and other inputs. In some policy circles, this literature is used to justify targeting agricultural interventions to women. If they are intrinsically as productive as men, then it can increase aggregate productivity to direct more opportunities and resources to women.

One of the earliest analyses of productivity differences considered maize farmers in Kenya. Moock (1976) found statistically insignificant differences between men and women plot managers, after accounting for input use, farmer characteristics and access to extension service. A number of similar studies have since focused on maize production. In Nigeria, Adeleke et al. (2008) compare the production functions for men and women maize farmers and find no significant difference in productivity after controlling for input use. Considering maize, beans and cowpeas in Kenya, Saito et al. (1994) use their analysis as a basis for a simulation exercise and conclude that if women had the same access to resources as men, output would increase by 22%.

In a similar analysis in Nigeria, Saito et al (1994) analyze the gender productivity gaps at both the household and the plot level. When they compare the productivity of male- and female-headed households, they do not find a gender gap after controlling for resource use, plot and farmers' characteristics. However, they do find that men plot managers are more productive than women plot managers, even conditional on input use.

Most of these analyses compare the productivity differences of men and women across the entire sample. In an approach to see if there are productivity differences between men and women in the same households, Udry (1996) in Burkina Faso found that these intra-household differences in productivity across men's and women's plots were large and could be attributed to differences in the intensity with which labor and fertilizer were applied to these plots. He estimated that shifting labor and fertilizer from men's plots to women's plots within the same household would substantially increase total household output. This is one of the few studies that considers the impact of reallocating inputs, rather than simply providing women with an amount similar to that of men.

In some regions of the Sahel, some plots are farmed individually by men or women, while others are farmed by the household. Guirkingner et al. (2015) find that in Mali, even after controlling for crop and household fixed effects, the value of the crops grown on women's individual plots and on collective family plots is lower than on men's individual plots. This contrasts somewhat with the findings of Kazianga and Wahhaj (2013) who don't find a gender gap between men's and women's individual plots but find instead that household plots, which are owned by the household head and farmed collectively, are more productive. And de la O Campos et al (2016) use a decomposition analysis and find that the productivity gaps are higher between male-only managed and female-only managed plots than between female-only managed plots and those jointly managed. These studies caution us against simply dividing plots into men's and women's and suggest that we may also need to consider joint plots and other management arrangements.

All of the results discussed thus far are for staple crops. Much less work explicitly considers cash crops, in part because the assumption is often made that women play a relatively small role in cash crop production. A recent study by Hill and Vigneri (2011) on cocoa farmers in Ghana, finds no gender gap in productivity after controlling for resource use (labor, fertilizer, insecticides, farm equipment; rainfall and farm quality).

Some of the more recent literature wrestles with the underlying methodological issues. Many of the studies above only consider one crop; one reason that women obtain lower values for their agricultural production overall may be due to the mix of crops that they produce. Even if men and women are equally productive in producing each crop, if women systematically choose to produce more of the lower value crops (or if households assign them to this role), their overall profitability will be lower.

One approach is to model both the crop choice and the level of output. For example, wa Githinji et al. (2014) use a two stage least squares estimation to test the productivity differentials between men and women farmers in Kenya.⁴ In the first stage, they estimate crop choice by regressing it on a set of control variables and instruments (exogenous factors affecting farmers' crop choice but not affecting productivity). In the second stage, they include the estimates of crop choice to assess the factors related to farm productivity. After accounting for endogenous crop choice in this way, they find no significant productivity differential between men and women farmers. This suggests that although the value of women's crops is lower than that of men, the disparity primarily reflects differences in the crops grown.

Similarly, Peterman et al. (2011) observe significant differences in major crops grown by men and women farmers in Nigeria. After controlling for inputs and for the characteristics of households and farmers, they still found significant gender differences in productivity. However, these differences are driven by which crops are grown, indicating primary crop choice as an important factor in determining gender gaps in productivity.⁵

These estimations assume that inputs are exogenously determined, rather than considering that the use of such inputs is a choice, albeit a constrained choice. It may not be appropriate to treat the inputs as exogenous, particularly when considering gender differences in productivity. For example, men and women may use different amounts of fertilizer because their other labor commitments differ, affecting the time available for applying fertilizer. Alternatively, their fertilizer use may reflect differential access to markets, or differential ability to mobilize cash or credit to make purchases. Thus, gender may condition whether or not one uses inputs.

The problem of endogenous input use is one that has given rise to a number of econometric approaches. One is the estimation of a restricted profit function; a dual approach whereby input demand and output supply can be derived from the profit function. In this approach, input demand functions are obtained, based on the assumption that profit is maximized. From the input demand functions, factor shares and the corresponding price elasticity of profit can be obtained. Using a restricted profit function allows the investigation not only of technical efficiency, which considers the level of output obtained from a given level of inputs, but also provides insights into allocative efficiency, which considers whether the most profitable combination of inputs is used by equating the marginal product of a variable with its price. The relative allocative efficiency of men and women can be compared; in addition, we can test whether farmers have absolute allocative efficiency, that is, whether the marginal product of each factor is equal to its price.

This approach has not been widely used for testing gender differences in productivity. Adesina and Djato (1997) employ a restricted profit function to measure technical and allocative efficiency of men

⁴ They also compare male- and female-headed households in Uganda.

⁵ In Uganda, there is significant gender gap in yield of sorghum and sweet potato, while for other crops it is insignificant. In Nigeria, the gap is driven by choice of tomatoes as primary crop.

and women rice farmers in in Cote d'Ivoire. While the fixed inputs (land and capital) and the prices of variable inputs (labor and fertilizer) are correlated with profit, the sex of the farmer is not. They find that the technical and allocative efficiency of men and women farmers is not different.

Overall, the literature applying production or profit functions finds little or no significant difference in male and female agricultural productivity or technical efficiency after controlling for access to inputs and for characteristics of plots, households and farmers. This body of evidence has frequently been used to support a claim that women are as productive as men in farming – and that any gender differences in productivity are due to disparities in access to resources. This evidence is used to support the claim that policy should increase women's access to resources. Although increasing the amount of resources available to farmers should increase overall productivity, the evidence doesn't necessarily suggest that the best way to allocate an additional unit of inputs is to provide it to women rather than men. A number of assumptions, many of which may be valid, would have to be made to support this claim.

Although the literature does point to the importance of obstacles that limit women's access to productive resources, this alone does not directly result in clear policy implications. Changing the constraints faced by women will require more insights into the local context to understand why there are gender differences in access to resources and how to resolve them.

Estimating Labor Productivity

A related set of studies compares the labor productivity of men and women in agriculture. This approach uses detailed information on the labor inputs to identify how they are related to output levels. The advantage of this approach is that it does not require that output be disaggregated into men's and women's output. Instead, it allow us to analyze how the labor inputs of men and women affect total farm productivity.

Measuring labor inputs presents many challenges, even before we begin to disaggregate by gender. Labor is often measured in days worked; farmers are asked to estimate how many days were spent working on each task. Often farmers are asked to recall this for the previous farming season. It implicitly assumes that a day's work is a useful measure of input and that the contribution of each day's work is equal. These measure rarely account for hours worked or effort expended.

Male and female labor may not be perfect substitutes. For a variety of reasons, including social norms, skills, and physical capabilities, there may be differences in the labor provided by men and women. In addition, because women typically have much higher burdens in terms of the time and physical requirements of providing care labor – such as caring for children, preparing meals, fetching water, and doing laundry – they have less time available for agricultural work. The opportunity costs of labor for men and women may also differ, since the non-farm opportunities for earning income may be lower for women.

A challenge in estimating the contributions of men's labor and women's labor separately is that the usual measure of labor input – time spent working on a farm or plot – does not necessarily account for the knowledge or management skills. A family member might work few hours but provide knowledgeable direction to others; this will not be well accounted for in the data and the results may be less meaningful.

Nevertheless, the usual approach is to treat men's labor and women's labor as separate inputs in an

agricultural production function. To determine the marginal contribution to output, the indirect estimates of production elasticities of different inputs can be derived. In simple terms, household output value (or some other measure of production) is regressed on a set of inputs, including men's labor and women's labor. The elasticities can be compared to one another; it is straightforward to test statistically the hypothesis that the two elasticities are identical.

In this literature, results can be quite sensitive to the specification of the functional form for the production function. As a result, it is desirable to use a relatively flexible functional form. Most studies that estimate the distinct labor contributions of men and women use a translog production function, which is a flexible form that nests many other common specifications. The translog is estimated from household-level data on output and inputs, including men's labor and women's labor. This estimation gives the marginal product of factor inputs, including labor.

Using this method for a study of rice farmers in Nepal, Aly and Shields (2010) find that there is no gender gap in labor productivity; the estimations account for differences in the quality of land and technology applied. Kumar and Hotchkiss (1988) use a similar approach in Nepal to study a range of crops and estimate it by season. They find that the marginal productivity of labor varies across seasons and crops. In the dry season, women's marginal product of labor is significantly higher than that of men for wheat, maize and mustard but is lower for paddy. In the wet season, women's marginal productivity of labor is lower than men's for most crops. They attribute these differences to women's other responsibilities across seasons.

In Peru, Jacoby (1992) finds that marginal product of female labor to farm output is significantly lower than that of male labor. He argues that this result can probably be attributed to the fact that women are casual workers in cropping, while their labor contribution is much higher in livestock production and household production. Laufer (1985) utilizes ICRISAT data for six Indian villages in 1975-76 and 1976-77 and found that women's marginal product of labor is lower for all three crops – sorghum, rice and legumes, in the analysis. Consistent with the gender division of labor, he finds that male and female labor are not necessarily substitutes; for some crops they are complements.

Because men and women typically specialize fully or partially in different types of farm labor, it is difficult to interpret these findings. If we know that on the margin, an hour of women's time spent weeding contributes less to overall productivity than an additional hour of men's time spent in preparing the land, how would this change the types of interventions that are proposed? If the marginal product of women's labor is high, this could imply that it would be useful to free up some additional time for women to contribute to agriculture. But it is equally true that if the marginal product of women's labor turns out to be low, it might support an argument that there are particularly high payoffs to redressing this imbalance.

Stochastic Frontier Analysis

Stochastic frontier analysis (SFA) is a parametric technique that allows for estimating the efficiency of individual production units relative to a frontier. These efficiency measures can then be regressed on a variety of explanatory variables to see how much of the efficiency differences can be accounted for by different variables.

In this approach, output is regressed on factors of production and then the estimates of inefficiency/error term are regressed on factors (e.g., farmers' characteristics, plot and household characteristics) hypothesized to result in output being lower than the maximum. This approach is widely used in the agricultural economics literature, but has only been used to a limited extent to consider whether inefficiency is correlated with sex of the farmer.

Kinkingninhoun-Médagbé et al. (2010) analyze data on rice farmers in Benin and find no gender gap in technical efficiency, after controlling for date of planting (related to access to ploughing equipment), irrigation access and experience in farming. This suggests that women are as efficient farmers as men; to the extent that their measured productivity is lower than that of men, it is explained by unequal access to irrigation facilities and ploughing equipment. In a small sample but long term study of coffee farmers in Papua New Guinea, Overfield and Fleming (2001) find that male and female labor do not have statistically different effects on technical efficiency.

In Nigeria, among cassava farmers, Timothy and Adeoti (2006) find that women have higher technical efficiency but men have higher allocative efficiency. However, both are operating below optimum levels of output. They suggest that women have lower access to resources and therefore, by shifting resources from men's plots to women's plots, overall efficiency can be increased.

Again, this literature can be helpful in identifying the constraints that farmers face in attaining higher levels of efficiency. But it doesn't provide clear guidance on how to effect these changes and what types of interventions would have the highest returns.

Decomposition of productivity differentials

It has been widely documented that men and women have different access to productive inputs for agriculture, including land, labor, seeds, fertilizer, and information. Decomposition analyses provide a means to identify how much of the productivity gap is due to this differential access and how much is due to the differential returns to these inputs. For example, the studies that have used this approach have typically shown that there is both a difference in the amount of fertilizer used by men and women farmers and a gender difference in the returns to using fertilizer.

The most frequently used technique is a Oaxaca-Blinder decomposition. This distinguishes between (i) the endowment effect or the differences in average characteristics of productivity generating factors, such as land, labor and other inputs; and (ii) the structural effects, which are gender differences in the returns to these factors. A recent set of papers use the LSMS-ISA data collected by the World Bank and Oaxaca-Blinder decomposition analysis to identify these patterns in six countries. These results are summarized in *Leveling the Field* (World Bank and ONE Campaign, 2014).⁶ This analysis uses the value of total crop output per hectare as the measure of productivity and compares plots managed by men with those managed by women. Analyses from five countries (Ethiopia, Malawi, Niger, southern Nigeria, and Uganda) find gender gaps in productivity when simply comparing the differences in value of output.

Controlling for plot size and geographic factors, the gender differences generally increase, rather than

⁶ The *Leveling the Field* report is based on many of the papers cited in this section. The comparative analysis here is drawn from the report. In addition, many of the individual papers are cited for their country-specific contributions.

decrease. In Niger, the gender gap increases from 19% to 66%. While it disappears in Nigeria (South), a gender gap of 46% results in Nigeria (North), and a gap of 33% is found in Uganda.

One of the reasons for these dramatic differences is related to farm size. For Malawi, Niger, Nigeria (North and South), and Uganda, after accounting for the differences in farm size, the gender gap widens. Women, on average, have smaller holdings than men. In these countries, among smallholder farms, there is an inverse relationship between farm size and productivity; smaller farmers are more productive per unit of land. Using this information to simulate what would happen if women had larger farm sizes, similar to those of men, they find that women would be even less productive. This finding has the perverse implication that the productivity of land could be increased by dividing farms and making them smaller. It underscores that a focus on land productivity obscures the fact that we may care much more about the productivity of labor, which is more closely linked to the well-being of the poor.

The main objective of these studies is to move beyond quantifying the gender gap to identify the contribution of each factor to the gender gap in productivity. The endowment effect is significant and in men's favor for Malawi and south Nigeria, while in north Nigeria, Uganda and Tanzania, it is in women's favor. In Ethiopia, both endowment and structural effects are insignificant. Significant positive structural effects are found in Malawi, Niger, northern Nigeria, Uganda and Tanzania, indicating that most of the gap is due to gender differences in returns to factors. The studies also perform detailed decompositions of gender gaps to identify the factors that contribute to the endowment and structural effects. For example, in Tanzania, plots being farther from a family's homestead is associated with a female structural disadvantage in productivity.

The individual country analyses provide more detail. In Uganda, because so few purchased inputs are used (fertilizer, chemicals, and improved seed) by any farmers, women's lack of access to these does not contribute much to the gender productivity gap (Ali, 2015). In Nigeria, there are significant differences in gender productivity in the North and in the South. In the north, not only do women have less access to inputs, but the returns to these factors of production are also lower (Oseni, 2015).

To understand the patterns of gender differences it is useful to have information not simply on a median household or an average household, but across the entire productivity distribution. The results overall suggest that the gender gap in productivity is highest above the middle of productivity distribution. For Ethiopia, access to resources matters more at the upper end of the productivity distribution (Aguilar et al, 2015). At lower end of the distribution, most of the gap is explained by endowment effect, which implies that at lower levels of productivity, men and women have similar returns to factors and residual gender gaps are driven by women's lower access to inputs. A different explanation is offered by Slavchevska (2015, p. 349), "at the lower end of the productivity distribution, women command lower returns to the factors of production, but their smaller plots mask the gender gap in returns." Moving up the productivity distribution, structural effects become more important than endowment effects. Distance to plot or distance to nearest market, returns to labor use and child dependency ratio are the consistent factors contributing to female structural disadvantage.

Technology adoption:

When detailed production data are not available, researchers have sometimes used technology adoption as a proxy for increased agricultural productivity. The implicit assumption is that farmers will not adopt

these improved technologies if they do not increase productivity, although a more appropriate assumption is that farmers will not adopt them unless they increase profit.

This approach does not require detailed information on the use of inputs and outputs. In particular, it does not require that the researchers calculate a value for labor inputs; instead, we simply can assume that farmers take the labor requirements into consideration when deciding whether or not to adopt. For example, research by Duflo, Kremer, and Robinson (2011) suggests that applying a half teaspoon of fertilizer as top dressing to maize plants is profitable in Kenya, but farmers have been reluctant to adopt this technology. This suggests, in part, that farmers place a higher value on their labor than do the researchers.

An extensive literature attempts to understand which farmers use improved technologies and why more farmers do not do so. In their simplest form, these studies use regression analyses to identify the characteristics of adopters. Including a variable for the sex of the farmer provides a simple way to see if the patterns differ for men and women. A number of studies suggest that, all else equal, men and women are equally likely to adopt these technologies. The challenge is that it is rare that all else is equal. Many technology adoption studies simply add a dummy variable for the sex of the household head although, as noted earlier, this tells us little about the differences between men and women.

Doss and Morris (2001) analyze the adoption of improved maize seed and fertilizer in Ghana. They find that holding other inputs constant, men and women farmers in male-headed households were indistinguishable in their adoption of these technologies. However, women farmers in female-headed households were less likely to use these improved technologies. This suggests that there are unobservable differences between women living in male- and female-headed households. It may be that women living in male-headed households have access to information and some unobservable resources that are unavailable to women living in households without men. But if we only consider the sex of the household head, we will miss the women farmers living in male-headed households who are adopting technologies.

Similar results were found by Chirwa (2005) for adoption of fertilizers and hybrid seed varieties by smallholders in Malawi. After controlling for plot, household level, institutional and farmers' characteristics, and infrastructure, the study found no significant gender differences in adoption of technology. Adding household headship to the analysis, however, suggests that women farmers in female-headed households are less likely to adopt fertilizers but equally likely as men and women in male-headed households to adopt hybrid seeds.

One recent study considers how an extension intervention affects the adoption of mineral fertilizers, row planting and improved legume varieties in the DRC (Lambrecht 2014). Conditional on plot, village and farmer's characteristics, the likelihood of adoption is increased when men and women participate jointly in the extension program.⁷ The extension program also had a positive impact on the adoption of mineral fertilizers and improved legumes by women in female-headed households. The likelihood of technology adoption is insignificant when the program targets only women in male-headed households, except for mineral fertilizers.

⁷ The joint participation includes men and women from the same household.

Two analyses consider the actual provision of inputs to men and women in Malawi, rather than just estimating what women's productivity would be if they used the same inputs as men. Karamba and Winters (2015) find that the provision of fertilizer subsidies at the household level increases productivity on both plots managed by men and those managed by women. But the increases in women's productivity are not disproportionately large enough to reduce the gender gap in productivity. An earlier study, found that when men and women farmers were provided with inputs as part of a national trial, there was not a gender gap in productivity (Gilbert et al, 2002).

What can we learn from looking at technology adoption? It is useful to know the extent to which men and women are using various technologies and the extent to which differential technology adoption can be explained by access to land, information, and other relevant inputs. But even if we find that there are unexplained gender differences in technology adoption, this doesn't tell us whether there are potentially new technologies that women farmers would adopt to increase their productivity.

Summary

A large literature has sought to compare the productivity of men and women in agricultural activities in developing countries. This proves to be a challenging task, starting at a conceptual level and continuing to measurement, estimation, and interpretation. Seen from a distance, however, it is not always clear what question is being addressed in this literature. What could we hope to find from this literature, even if all the conceptual and measurement issues could be solved, and the endogeneity problems made to disappear? If we find that women have lower measured productivity than men, what can we conclude about the potential returns to investments that target women? If we want to increase aggregate productivity, should we invest less in women, since they are less productive? Or more, since there would seem to be easier opportunities to increase their productivity with off-the-shelf technologies and interventions? If we find that productivity differs, but only because of inputs, what is an appropriate response? The literature has little to say to these questions; even the most powerful econometric tools do not allow us to answer the key policy questions.

5. Measuring Women's Productivity within Joint Production Systems

This section addresses a further difficulty that arises in estimating women's productivity in agriculture. Most of the academic literature assumes that men and women have distinct responsibilities and possibly distinct plots.⁸ But in reality, most farming takes place in households where men and women share responsibility for the same plots and for the whole of the household enterprise. It may be the case that there is a gender division of labor in these household enterprises, with men and women focusing on different agricultural tasks. And it may be the case that the management is not shared equally. Nevertheless, both men and women are involved in much of the agricultural production, and their contributions are not easily separated.

It is not a coincidence that all of the studies comparing the productivity on men's and women's plots were done in Africa. It is much less common in Asia and Latin America for women to manage separate plots. Even when women do manage distinct plots, these are often viewed as kitchen gardens or

⁸ The lone exception is the literature discussed above in the section on stochastic frontier analysis that seeks to estimate the marginal productivity of men's and women's labor to household output or profits.

homestead plots, rather than agricultural plots. And even in some places in Africa, there are few “women’s plots” other than those managed by women who are the heads of their households. Wives may not be reported as farmers with their own plots.

Thus, if we want to consider men’s and women’s contributions to agricultural productivity – and their potential contributions to agricultural productivity – we must find ways to consider women’s contributions to production activities that are jointly managed with men – or perhaps entirely managed by men.

The studies focusing on the marginal product of men’s and women’s labor provide some information about whether labor is being allocated efficiently within a household, especially when the labor is not perfectly substitutable. Yet, these studies don’t tell us anything about women’s contributions to farm management or decision-making. In joint farming systems, both men and women may be involved in management; the degree of involvement may differ and they may be responsible for different sets of decisions.

One option would be to do a much better job of identifying who is involved with each plot of land and finding ways to characterize farm labor that incorporates a broader range of management approaches than simply that it was done by a man or a woman. A useful beginning point would be to distinguish the productivity on plots that were jointly managed from those that were individually managed.

In principle, one way to test these questions would be to design a randomized controlled trial (RCT), introducing a new agricultural technology and then randomly providing it to men and women farmers separately and in settings where they farm jointly. The outcome of interest would be to see the differential ways in which the intervention alters household allocations depending on how the technology is targeted. But even such an experiment would not actually answer the key policy question. Such an experiment simply asks, for a given technology, does it matter how it is distributed? But it cannot address the bigger question of which interventions, targeted at whom, will have the highest social returns.

The key issue is whether it makes sense to target investments to women on their own, or women within joint systems, or perhaps not to target women at all. It may be important to target women, not necessarily because they are farming their own plots, but because they are deeply involved in providing labor and management to activities within the broader household production system. We have numerous examples suggesting that when new technologies are introduced that require women’s labor but where women receive few or none of the benefits, the technologies are not widely adopted. And there is preliminary evidence beginning to accumulate that providing women with information or training may increase household production, even when men are considered the primary farmer or manager.⁹

Thus, we have to consider two issues – how does targeting women and providing them with additional resources, technologies, or information affect the total productivity of the household? And how do we ensure that women benefit from the additional output and are thus willing to contribute? This second issue has serious implications for program design. It may involve including payments to multiple people within the same household in commercialization schemes or designing contracts so that women’s roles and contributions are visible and compensated.

⁹In an analysis of a dairy development project in Mozambique, Johnson et al (2015) find that training two people within the household, instead of only the male household head, resulted in higher levels of milk production.

Simply knowing who contributes the most at the initial stage is useful when designing appropriate interventions, but is not sufficient information on existing gender relations. Descriptions of the baseline situation also will not necessarily tell us where the potential exists for transforming gender relations and increasing productivity so that everyone benefits.

The literature on joint production raises some interesting possible directions for policy – or at least for further research. For example, there is consensus that secure property rights promote investment in agriculture; those who are worried about losing their land are less likely to invest. In one of the few analyses that consider gender, Goldstein and Udry (2008) examine land rights in Ghana and find that the women invest less (the investment that they consider is fallowing the land and allowing the fertility to recover) due to their weaker land tenure rights which is a result of their lower political power within the community. In this case, they are considering an area where men and women farm separate plots. But no research has specifically considered the extent to which women may be deterred from investing in agriculture when they may lose access to their land if their husband dies or divorces them. The emerging literature on joint titling of land may be able to provide some insights into this issue.

6. Conclusions:

What do we know about women and agricultural productivity in developing countries? How have new data sources shaped and improved our understanding? Will interventions that specifically target women lead to improvements in aggregate agricultural productivity? Will female-focused agricultural interventions generate higher rates of return than interventions that are, in some sense, gender-blind?

The answers to these questions are complex. The emergence of new data sources has allowed for a greatly enriched understanding of gendered patterns of technology adoption, labor use, and productivity in agriculture. New data has rejected some of the groundless and unrealistic claims that were made in the past about the share of women's labor in agriculture, or the share of land owned by women; these pseudo-statistics continue to live on, however, in the echo chambers of the advocacy world.

At the same time, the data show in rich detail that women are deeply involved in all phases of agricultural production. Their labor at the plot level accounts for about 40 percent of the total field work in crop agriculture in Africa (Palacio-Lopez, 2015) – and it should be noted that we still lack data on women's participation in other areas of farming, including seed selection and management, input purchasing, output marketing, processing, and animal care, not to mention home activities such as food preparation and cooking. There is a widespread consensus that women devote more time than men to many of these activities, and certainly spend more time than men on household labor. As economies continue to grow and transform, we are also reminded that agriculture is only one sector. In some settings, men may be moving out of agricultural work into other sectors that provide better opportunities. This may leave women with greater responsibility for agricultural production, or at the very least it may alter the basis on which joint activities are organized.

The new data have allowed for further and methodologically more sophisticated analyses of the relative productivity of men and women. These analyses face a host of conceptual and methodological challenges, but most studies suggest that women would be approximately as productive as men as farmers, given access to similar resources. Some studies have been interpreted to suggest that there is a potential gain in overall productivity by targeting women, but this evidence is not clear, and there are

also studies that suggest the opposite. As the universe of available data expands, it may be possible to identify the conditions under which women have higher productivity and which factors systematically constrain women's output.

Some subjects continue to be poorly addressed in survey data, such as women's land tenure security within households and communities; i.e., women's ability to retain land in the event of divorce or the death of a spouse. The literature also draws heavily from research from Africa, comparing plots managed by men with those managed by women. The rest of the world may look very different. In particular, in much of South Asia, the poor are landless, and the challenges for women in this context may look very different from those of African women who are farming their own plots of land.

Nevertheless, the emerging data can provide important baseline information as we design projects and interventions. They reveal the strongly gendered patterns of labor use and household responsibilities, and in this way they can help identify the constraints that limit men's and women's productivity.

The new data and the resulting analysis do not, however, answer all questions. In particular, they do not tell us very directly how to target interventions. If we find that women provide 40% of the labor in crop agriculture, what does this fact tell us about whether (and how) to target interventions to women? Should we follow a simple rule of thumb and provide 40% of the resources to women and 60% to men? Should we ignore women altogether since their labor input is less than that of men? Such an approach would seem to ignore the more important question – namely, where can we expect to find the highest rates of return.

Arguably, some of the narrow analyses of men's and women's productivity in agriculture have taken us away from the bigger and more important questions. The academic search for more precise estimates has spawned a rich and challenging literature, but it has diverged from several of the most important policy questions.

Can we target women for agricultural development interventions? If so, how? There are a number of promising approaches for targeting women in the agricultural sector.¹⁰ Some of these interventions are targeted specifically at women while others seek ways to ensure that women are able to participate in projects that do not have a gender focus. It is important to note that key interventions for increasing women's agricultural productivity might not lie in the agricultural sector at all. Any interventions that reduced women's drudgery – such as by providing tap water or reducing the daily trek for firewood, could free up time and energy which could be productively used in agriculture. Improving child health and nutrition could also reduce the care burdens that women face and could free up women's time for higher productivity activities in agriculture. But these interventions are hardly ever evaluated in terms of their impact on women's agricultural productivity.

In addition, this focus on comparing men's and women's agricultural productivity may lead us away from asking how interventions in the agriculture sector that include women may have impacts not only on agricultural output, but also on reducing poverty and increasing gender equity.

Whether or not specifically interventions should explicitly *target* women rather than men, it is clear that a *gender-blind* approach to designing interventions will miss out on key constraints, opportunities, and impacts. Gender is embedded in the distribution of essentially all the resources used in agriculture – including land ownership, farm management decisions, market access for inputs and outputs, information from extension services, use of ICT, etc. If interventions fail to consider how gender is

¹⁰ See Quisumbing and Pandolfelli (2010) and Doss et al. (2012) for reviews of such projects.

embedded in the system, we will miss critical opportunities for transforming agricultural systems , increasing productivity, reducing poverty and improving people’s lives.

References

- Adeleke, O A, O I Adesiyun, O A Olaniyi, K O Adelalu, and H M Matanmi. 2008. “Gender Differentials in the Productivity of Cereal Crop Farmers: A Case Study of Maize Farmers in Oluyole Local Government Area of Oyo State.” *Agricultural Journal* 3 (3): 193–98.
- Adesina, Akinwumi A, and Kouakou K Djato. 1997. “Relative Efficiency of Women as Farm Managers: Profit Function Analysis in Cote d’Ivoire.” *Agricultural Economics* 16: 47–53.
- Aguilar, Arturo, Eliana Carranza, Markus Goldstein, Talip Kilic, and Gbemisola Oseni. 2014. “Decomposition of Gender Differentials in Agricultural Productivity in Ethiopia.” Policy Research Working Paper No. 6764. Washington D.C.: The World Bank.
- Ali, Daniel, Derick Bowen, Klaus Deininger, and Marguerite Duponchel. 2015. “Investigating the Gender Gap in Agricultural Productivity: Evidence from Uganda.” Policy Research Working Paper No. 7262. Washington D.C.: The World Bank.
- Aly, Hassan Y, and Michael P Shields. 2010. “Gender and Agricultural Productivity in a Surplus Labor, Traditional Economy: Empirical Evidence from Nepal.” *The Journal of Developing Areas* 43 (2): 111–24.
- Backiny-Yetna, Prospere, and Kevin McGee. 2015. “Gender Differentials and Agricultural Productivity in Niger.” Policy Research Working Paper No. 7199. Washington D.C.: The World Bank.
- Bardasi, Elena, Kathleen Beegle, Andrew Dillon, and Pieter Serneels. 2011. “Do Labor Statistics Depend on How and to Whom the Questions Are Asked? Results from a Survey Experiment in Tanzania.” *World Bank Economic Review* 25 (3): 418–47.
- Chirwa, Ephraim W. 2005. “Adoption of Fertiliser and Hybrid Seeds by Smallholder Maize Farmers in Southern Malawi.” *Development Southern Africa* 22 (1): 1–12.
- de la O Campos, Ana Paula, Katia Alejandra Covarrubias and Alberto Prieto Patron. 2016. “How Does the Choice of the Gender Indicator Affect the Analysis of Gender Differences in Agricultural Productivity? Evidence from Uganda.” *World Development* 77: 17-33.

- Doss, Cheryl. 2014. "If Women Hold up Half the Sky, How Much of the World's Food Do They Produce?" In *Gender in Agriculture: Closing the Knowledge Gap*, edited by Agnes R Quisumbing, Ruth Meinzen-Dick, Terri L Raney, André Croppenstedt, Julia A Behrman, and Amber Peterman, 69–88. Dordrecht: Springer. Chapter 4.
- Doss, Cheryl, Chiara Kovarik, Amber Peterman, Agnes Quisumbing, and Mara van den Bold. (2015). "Gender inequalities in ownership and control of land in Africa: myth and reality." *Agricultural Economics* 46(3): 403-434.
- Doss, Cheryl, and Michael L Morris. 2001. "How Does Gender Affect the Adoption of Agricultural Innovations? The Case of Improved Maize Technology in Ghana." *Agricultural Economics* 25: 27–39.
- Doss, Cheryl R, Zoe Bockius-Suwyn, and Shereen D'Souza. 2012. "Women's Economic Empowerment in Agriculture : Supporting Women Farmers." Washington D.C.: UN Foundation.
- Duflo, Esther, Michael Kremer, and Jonathan Robinson. 2011. "Nudging Farmers to Use Fertilizer: Theory and Experimental Evidence from Kenya." *American Economic Review* 101 (6): 2350–90.
- Fisher, Monica, and Vongai Kandiwa. 2014. "Can Agricultural Input Subsidies Reduce the Gender Gap in Modern Maize Adoption? Evidence from Malawi." *Food Policy* 45: 101–11.
- Gilbert, Robert. A. Sakala, Webster D., & Todd D. Benson. 2002. "Gender analysis of a nationwide cropping system trial survey in Malawi." *African Studies Quarterly*, 6(1-2), 223–243.
- Goldstein, Markus, and Christopher Udry. 2008. "The Profits of Power: Land Rights and Agricultural Investments in Ghana." *Journal of Political Economy* 116 (6): 981–1022.
- Guirkinger, Catherine, Jean-Philippe Platteau, and Tatiana Goetghebuer. 2015. "Productive Inefficiency in Extended Agricultural Households: Evidence from Mali." *Journal of Development Economics* 116: 17–27.
- Hill, Ruth Vargas, and Marcella Vigneri. 2011. "Mainstreaming Gender Sensitivity in Cash Crop Market Supply Chains." ESA Working Paper No. 11-08. Rome: FAO.
- Jacoby, Hanan G. 1992. "Productivity of Men and Women and the Sexual Division of Labor in Peasant Agriculture of the Peruvian Sierra." *Journal of Development Economics* 37: 265–87.
- Johnson, Nancy, Jemimah Njuki, Elizabeth Waithanji, Marinho Nhambeto, Martha Rogers & Elizabeth Hutchinson Kruger. 2015. "The gendered impacts of agricultural asset transfer projects: Lessons from the Manica Smallholder Dairy Development Program." *Gender, Technology and Development*, 19(2), 145–180.
- Karamba, Wendy and Paul C. Winters. 2015. "Gender and Agricultural Productivity: Implications of the Farm Input Subsidy Program in Malawi." *Agricultural Economics* 46: 357-74.

- Kazianga, Harounan, and Zaki Wahhaj. 2013. "Gender, Social Norms and Household Production in Burkina Faso." *Economic Development and Cultural Change* 61 (3): 539–76.
- Kieran Caitlin, Kathryn Sproule, Cheryl Doss, Agnes Quisumbing, and Sung Mi Kim. Examining gender inequalities in land rights indicators in Asia. *Agricultural Economics* 46. <http://doi.wiley.com/10.1111/agec.12202> rg
- Kilic, Talip, Amparo Palacios-Lopez, and Markus Goldstein. 2015. "Caught in a Productivity Trap: A Distributional Perspective on Gender Differences in Malawian Agriculture." *World Development* 70: 416–63.
- Kinkinginhoun-Médagbé, Florent M, Aliou Diagne, Franklin Simtowe, Afiavi R Agboh-Noameshie, and Patrice Y Adégbola. 2010. "Gender Discrimination and Its Impact on Income, Productivity, and Technical Efficiency: Evidence from Benin." *Agriculture and Human Values* 27 (1): 57–69.
- Kumar, Shubh K, and David Hotchkiss. 1988. "Consequences of Deforestation for Women's Time Allocation, Agricultural Production, and Nutrition in Hill Areas of Nepal." Research Report No. 69. Washington D.C.: International Food Policy Research Institute.
- Lambrecht, Isabel. 2014. "Adoption of Integrated Soil Fertility Management in Burundi and Eastern DR Congo." PhD diss., University of Leuven. Leuven, Belgium
- Laufer, Leslie A. 1985. "The Substitution between Male and Female Labor in Rural Indian Agricultural Production." Center Discussion Paper No. 472. New Haven, Connecticut: Economic Growth Center, Yale University
- Lilija, Nia, Thomas F. Randolph, and Abrahmane Diallo. 1998. "Estimating Gender Differences in Agricultural Productivity: Biases Due to Omission of Gender-Influenced Variables and Endogeneity of Regressors." Paper submitted to American Agricultural Economics Association Annual Meeting, Salt Lake City, Utah.
- Mooock, Peter R. 1976. "The Efficiency of Women as Farm Managers: Kenya." *American Journal of Agricultural Economics* 58 (5): 831–35.
- Oseni, Gbemisola, Paul Corralb, Markus Goldsteinc, and Paul Winters. 2015. "Explaining Gender Differentials in Agricultural Production in Nigeria." *Agricultural Economics* 46: 285–310.
- Overfield, Duncan, and Euan Fleming. 2001. "A Note on the Influence of Gender Relations on the Technical Efficiency of Smallholder Coffee Production in Papua New Guinea." *Journal of Agricultural Economics* 52 (1): 153–56.
- Palacios-Lopez, Amparo, Luc Christiaensen, and Talip Kilic. 2015. "How Much of the Labor in African Agriculture Is Provided by Women?" Policy Research Working Paper No. 7282. Washington D.C.: The World Bank.

- Peterman, Amber, Agnes R Quisumbing, Julia A Behrman, and Ephraim Nkonya. 2011. "Understanding Gender Differences in Agricultural Productivity in Uganda and Nigeria." *Journal of Development Studies* 47 (10): 1482–1509.
- Quisumbing, Agnes R., and Lauren Pandolfelli. 2010. "Promising Approaches to Address the Needs of Poor Female Farmers: Resources, Constraints, and Interventions." *World Development* 38 (4): 581–92.
- Saito, Katrine A., Hailu Mekonnen, and Daphne Spurling. 1994. "Raising the Productivity of Women Farmers in Sub-Saharan Africa." World Bank Discussion Paper No. 230. Washington D.C.: The World Bank.
- Slavchevska, Vanya. 2015. "Gender Differences in Agricultural Productivity: The Case of Tanzania." *Agricultural Economics* 46: 335–55.
- Timothy, Awoyemi Taiwo, and Adetola I. Adeoti. 2006. "Gender Inequalities and Economic Efficiency: New Evidence from Cassava-Based Farm Holdings in Rural South-Western Nigeria." *African Development Review* 18 (3): 428–43.
- Twyman, Jennifer, María Alejandra García, and Juliana Muriel Muriel. 2015. "Identifying Women Farmers: Informal Gender Norms as Institutional Barriers to Recognizing Women's Contributions to Agriculture." *Journal of Gender, Agriculture and Food Security*, 1(2), 1–17.
- Udry, Christopher. 1996. "Gender, Agricultural Production and Theory of the Household." *Journal of Political Economy* 104 (5): 1010–1046.
- wa Githinji, Mwangi, Charalampos Konstantinidis, and Andrew Barenberg. 2014. "Small and Productive : Kenyan Women and Crop Choice." *Feminist Economics* 20 (1): 101–29.
- World Bank and the ONE Campaign. 2014. "Leveling the Field: Improving Opportunities for Women Farmers in Africa." Washington, DC.