www.elsevier.com/locate/worlddev

World Development Vol. 87, pp. 152–170, 2016

http://dx.doi.org/10.1016/j.worlddev.2016.06.006

Investigating the Gender Gap in Agricultural Productivity: Evidence from Uganda

DANIEL ALI, DERICK BOWEN, KLAUS DEININGER and MARGUERITE DUPONCHEL *

The World Bank, Washington, DC, USA

Summary. — Women comprise 50% of the agricultural labor force in Sub-Saharan Africa, but manage plots that are reportedly on average 20–30% less productive. As a source of income inequality and aggregate productivity loss, the country-specific magnitude and drivers of this gender gap are of great interest. Using national data from the Uganda National Panel Survey for 2009–10 and 2010–11 that include a full agricultural module and plot-level gender indicator, the gap before controlling for endowments was estimated to be 17.5%. Panel data methods were combined with an Oaxaca decomposition to investigate the gender differences in resource endowment and return to endowment driving this gap. Although men have greater access to inputs, input use is so low and inverse returns to plot size so strong in Uganda that smaller female-managed plots have a net endowment advantage of 12.9%, revealing a larger unexplained difference in return to endowments of 30.4%. One-half of this is attributed to differential returns to the child dependency ratio, implying that greater child care responsibility is the largest driver of the gap. Smaller drivers include differential uptake of cash crops, differential uptake and return to improved seeds and pesticides, and differential returns to male-owned assets.

Key words — gender, agriculture, productivity, decomposition, Sub-Saharan Africa, Uganda

1. INTRODUCTION

Women comprise 50% of the agricultural labor force in Sub-Saharan Africa, but manage plots that are 20–30% less productive than male-managed plots (FAO, 2011). This agricultural productivity gap contributes to income inequality between women and men. In some cases, the productivity gap is also partially driven by an inefficient over-allocation of inputs to male-managed plots, resulting in large aggregate productivity losses for the agricultural sector (Akresh, 2008; Udry, 1996). [1]

As sources of inequality and aggregate productivity loss, the country-level magnitude and drivers of the agricultural productivity gender gap (henceforth referred to as "the gap") are of great interest to policy makers. Investigation of the gap typically takes one of two approaches: (1) Inter-household analysis that uses agricultural production in female- and male-headed households as proxies for agricultural production on female- and male-managed plots, and (2) Intra-household studies that use production on female- and male-managed plots located within the same household as proxies for female- and male-managed production in the overall sector.

Inter-household analysis commonly estimates the simple gap in mean yield between male and female households, and then tests whether the gap is driven by differences in resource endowment (distribution of resources) or return to resource endowment (technical efficiency). [2] The most common procedure for this test is to run a multivariate regression on the pooled sample of female- and male-headed households and to observe whether there is a statistically significant correlation between gender of household head and yield after accounting for all observable endowments. If so, this remaining correlation is considered to result from a difference in technical efficiency between male- and female-headed households. Another possibility is to estimate the technical efficiency of male- and female-headed households using a stochastic production frontier and observe whether the mean difference in technical efficiency is statistically significant (Kinkingninhoun-Medagbe, Diagne, Simtowe, Agboh-Noameshie, & Adegbola, 2010; Oladeebo & Fajuyigbe, 2007).

Some inter-household analyses also test whether resources are allocated efficiently between male- and female-headed households (allocative efficiency). One approach is to perform the above multivariate regression analysis with individual inputs, such as fertilizer or labor, as the dependent variable in place of yield. If there is a statistically significant correlation between input use and gender of household head after accounting for all observable characteristics, the input is considered to be allocated inefficiently between households. That is, given declining marginal returns to the input in question, redistribution of the input from male- to female-headed households with the same characteristics but lower endowment of the resource in question would increase overall agricultural yield in the sector (e.g., Horrell & Krishnan, 2007). Another approach to test for allocative inefficiency is to estimate each input's marginal value product separately for male and female farmers and to observe whether this is above or below the input's factor price (Tiruneh, Tesfaye, Mwangi, & Verkuijl, 2001).

Inter-household analyses of the gender gap tend to find a statistically significant and practically substantial gap in mean yield (or value of yield) of around 20–30% in favor of male farmers. The overwhelming conclusion from these studies has been that the gap is driven by differences in resource endowment between male- and female-headed households, rather than technical efficiency (Horrell & Krishnan, 2007; Kinkingninhoun-Medagbe *et al.*, 2010; Oladeebo & Fajuyigbe, 2007), with a few exceptions (Holden, Shiferaw, & Pender, 2001; Quisumbing, 2001). Characteristics accounted for in the tests for technical efficiency include physical inputs (organic and chemical fertilizer, crop protection chemicals, improved seed varieties, mechanization), human capital (labor, education, extension services, childbearing), quality

* Final revision accepted: June 12, 2016

of land and water access, land investment (duration left fallow, erosion control and water harvesting structures, planting of tree crops), access to credit and input and output markets, and informal institutional constraints (norms regarding division of crop production and assignment of household responsibilities).

One limitation of these tests for technical inefficiency is that they do not quantify the individual contribution of endowment differences relative to the simple difference in agricultural productivity. A detailed decomposition of the endowment gap would be useful for designing policy programs to close the gap and highlighting areas for future in-depth study. Another limitation of inter-household studies is that, while the gap between male- and female-headed households is policy relevant in its own right, it is important to distinguish the inter-household gender gap from the overall gender gap. The majority of female-managed plots in Sub-Saharan Africa are located within male-headed households, the typical structure of which is fundamentally different from the typical structure of female-headed households. In particular, female-headed households are overwhelmingly characterized by cases where the husband has passed away, is a migrant laborer, or is polygamously married and member of a different household. Analysis of the inter-household gap therefore draws inference from a subset of female-managed plots quite distinct from the typical female-managed plot, and categorizes female-managed plots as "male" and vice versa. Finally, given fundamental differences between female-headed and male-headed households, these tests for technical inefficiency are likely to be invalid due to omission of important unobservable characteristics from the analysis.

To address the latter two concerns, intra-household analyses of the gender gap use plot-manager-level agricultural data from panel surveys and restrict the study sample to households in which both male- and female-managed plots are present. This allows these studies to use a fixed effects estimator to account for all unobserved household characteristics, as well as observed farmer and plot-level characteristics, substantially reducing the likelihood of omitted variables confounding the analysis. The most rigorous of these studies first show that unobserved household characteristics and a set of plot-level characteristics are insufficient to completely explain the gap between female and male productivity. As a result, they either introduce an omitted variable to explain the remaining gap, such as duration left fallow (Goldstein & Udry, 2008) or physical and labor inputs (Udry, 1996), or explore how the gap differs by household-level characteristics, either by introducing interaction terms (Akresh, 2008) or re-running the analysis on separate samples (Akresh, 2008; Peterman, Quisumbing, Behrman, & Nkonya, 2011).

Until recently, intra-household studies have relied on small samples within geographically limited settings that raise concerns about the external validity of results within or across countries. The seminal paper, Udry, Hoddinott, Alderman, and Haddad (1995), examines a sample of 150 households in six villages, while (Goldstein & Udry, 2008) examine 60 married couples in four enumeration areas. Only one study (Akresh, 2008) examined nationally representative data. The recent addition of gender disaggregated plot-level data to nationally representative panel surveys by the Living Standards Measurement Study Integrated Surveys on Agriculture (LSMS-ISA) has made it feasible to apply intra-household analysis to larger, more informative datasets.

The early studies that use intra-household analysis suffer from one of the same limitations of inter-household analysis in that, while they rigorously identify one subset of characteristics that contribute to the gender gap, they do not quantify the contribution of the complete set of individual endowments to the gap. In cases where the gap is not fully explained (Peterman *et al.*, 2011), the apparent difference in returns to endowments is not systematically explored.

A common solution to this problem in the labor economics literature has been to estimate the contribution of characteristics to the gap by examining changes in the coefficient on the gender indicator (the size of the gap) as covariates (characteristics) are sequentially introduced to the analysis. Gelbach (in press), however, shows that the change in the coefficient on the gender indicator depends on the order in which the covariates are introduced and that estimates from this procedure are, in that sense, path dependent. Sequential addition of covariates also does not quantify differences in returns to inputs and characteristics between male and female farmers. In some cases, returns to characteristics are estimated over the female and male samples separately and compared to more clearly identify areas where differences in returns exist (Hill & Vigneri, 2011). The relevance of these differences in returns is not clear, however, unless they are scaled to the size of the endowment for that characteristic.

An alternative procedure often applied in labor economics is the well-known Oaxaca–Blinder (O–B) Decomposition (Blinder, 1973; Oaxaca, 1973), described in detail in Section 4. Kilic, Palacios-Lopez, and Goldstein (2015) apply this procedure for the first time to the agricultural productivity gap. O–B decomposition is not path-dependent and, within a partial-equilibrium framework, quantifies the relative contribution of factors to the gap. It decomposes contributions to the simple difference into a component accounted for by endowments (endowment effect) and returns to these endowments (structural effect).

Kilic *et al.* (2015) find that female-managed plots in Malawi are 25% less productive than male-managed plots, and that 82% of this difference is explained by observable characteristics. They find that the primary contributions to the endowment gap are higher levels of adult male labor and selection of export crops, and that the primary contributions to the structural gap are the child dependency ratio (less time to devote to productive activities) and returns to male labor (possibly difficulty supervising male household labor). They use recentered influence function decomposition to show that these results hold throughout the productivity distribution.

Several studies have since replicated this procedure on cross-sectional datasets, including Aguilar, Carranza, Goldstein, Kilic, and Oseni (2015), Oseni, Corral, Goldstein, and Winters (2015) and de la O Campos, Covarrubias, and Patron (2016). These analyses are performed on cross-sectional data, however, that do not include within-household variation in gender of the plot manager. In that sense, the decomposition is still an arbitrary procedure in that it relies on the assumption that the mean of unobservable characteristics conditional on observable characteristics is equal to zero—the zero conditional mean assumption (Fortin, Lemieux, & Firpo, 2010)—and it is not clear which characteristics must be included in the analysis for this assumption to hold. To increase the likelihood that zero conditional mean holds, we introduce panel-data methods to the O–B decomposition of the gender gap in agricultural productivity to account for unobserved community, household and farmer characteristics. To our knowledge, Slavchevska (2015) is the only other study to combine Oaxaca Decomposition with household fixed effects.

A potential drawback of this approach is that it cannot separately account for the contribution of observable