ICAE 2024
32nd International Conference of Agricultural Economists
2-7 August 2024 I New Delhi I India

# Women's Work and Agricultural Productivity Gaps in India

Kajal Gulati[1], Koustuv Saha[1], and Travis J. Lybbert[3]

*1: Purdue University, IN, USA. 2: University of California – Davis, CA, USA*

*Corresponding author email: gulati28@purdue.edu*

**Abstract**

Most studies on gender gaps in agricultural productivity leverage within-household differences between plots managed by women and men. Such a gender-based division of plot management simplifies empirical tests for productivity differences, but it is not a common arrangement for agricultural households outside some locations in sub-Saharan Africa. In most rural households, women and men jointly participate in production, which complicates identification of gender-based productivity differences. This study proposes a broader empirical test of productivity gaps that applies to such systems, and that is rooted not explicitly in gender but in gender-based inequities. Specifically, we explore productivity gaps in rice-cultivating Indian households, where women and men perform specific and distinct cultivation tasks. We measure productivity gaps based on the differential use of family and hired female labor across households, then compare them with gaps based on the differential use of family and hired male labor. Using plot-level data, we identify significant gender-based productivity gaps after controlling for input use, plot- and household-level characteristics, and using village fixed effects and machine learning estimators to address selection and model misspecification concerns. Based on this identification strategy, households using family female labor have lower agricultural productivity, on average, than those also hiring female workers, such that foregone production value is greater than the cost of hiring women. We find suggestive evidence that this gap stems from skill differences between hired and family female workers. In contrast, we find no evidence of a similar gap based on the differential use of family and hired male labor. Overall, household welfare is lower because of gender inequities that shape women's work opportunities. These findings highlight the potential productivity implications of expanding women's labor choices, including both on- and off-farm job opportunities.

**JEL Codes:** D13, J16, J43, Q12

# Women's Work and Agricultural Productivity Gaps in India

Kajal Gulati[a], Koustuv Saha[a], and Travis J. Lybbert[b]

[a]Purdue University, West Lafayette, IN, USA
[b]University of California, Davis, CA, USA

**Abstract**

Most studies on gender gaps in agricultural productivity leverage within-household differences between plots managed by women and men. Such a gender-based division of plot management simplifies empirical tests for productivity differences, but it is not a common arrangement for agricultural households outside some locations in sub-Saharan Africa. In most rural households, women and men jointly participate in production, which complicates identification of gender-based productivity differences. This study proposes a broader empirical test of productivity gaps that applies to such systems, and that is rooted not explicitly in gender but in gender-based inequities. Specifically, we explore productivity gaps in rice-cultivating Indian households, where women and men perform specific and distinct cultivation tasks. We measure productivity gaps based on the differential use of family and hired female labor across households, then compare them with gaps based on the differential use of family and hired male labor. Using plot-level data, we identify significant gender-based productivity gaps after controlling for input use, plot- and household-level characteristics, and using village fixed effects and machine learning estimators to address selection and model misspecification concerns. Based on this identification strategy, households using family female labor have lower agricultural productivity, on average, than those also hiring female workers, such that foregone production value is greater than the cost of hiring women. We find suggestive evidence that this gap stems from skill differences between hired and family female workers. In contrast, we find no evidence of a similar gap based on the differential use of family and hired male labor. Overall, household welfare is lower because of gender inequities that shape women's work opportunities. These findings highlight the potential productivity implications of expanding women's labor choices, including both on- and off-farm job opportunities.

*Keywords:* Gender, agriculture, labor, productivity, women, India
*JEL Codes:* D13, J16, J43, Q12

---

This work was supported by the Bill and Melinda Gates Foundation and the U.S. Agency for International Development through the Cereal Systems Initiative for South Asia. The authors thank participants at the 2023 CGIAR Annual Gender Research and the 2021 Agricultural and Applied Economics Association Conferences for their comments; and David Spielman, Patrick Ward, Sumedha Minocha, Lorena Danessi, and Vartika Singh for their valuable contributions and insights. Any and all errors are the sole responsibility of the authors. This study was approved by the International Food Policy Research Institute Institutional Review Board (IRB #00007490; FWA #00005121). Corresponding author e-mail: gulati28@purdue.edu

# 1 Introduction

In most agricultural households globally, men and women jointly contribute to farm operations with the common objective of optimizing production and profits (Doss, 2018). While sharing responsibility for the cultivation of the same plots, they assume distinct roles and responsibilities, often organized by gender and shaped by prevalent gender norms (Doss, 2018; Quisumbing, 1996; Alesina, Giuliano, and Nunn, 2013). The multifaceted and interdependent nature of the on-farm tasks that men and women perform complicates disentangling their separate contributions to overall household production. Adding to this complexity is the patriarchal structure inherent in many of these joint production systems, where men manage most of the production operations. This not only makes it challenging to estimate the contributions of women in such production systems, but also opens the possibility of gender-based gaps in agricultural productivity and opportunities to improve productivity by addressing the source of these gaps.

In such joint agricultural production systems, we know little about the needs of women farmers, the constraints to their productivity, and ways to best support them, especially relative to their male counterparts (Doss, 2018; Quisumbing, 1996). Much of the literature on the constraints women farmers face and ways of expanding their productivity is limited to unique locations in sub-Saharan Africa where men and women within the same household manage separate plots (Quisumbing and Pandolfelli, 2010; Quisumbing, 1996; Udry, 1996; Akresh, 2005; Goldstein and Udry, 2008; Peterman et al., 2011; Kazianga and Wahhaj, 2013; Mahajan, 2019). Most of these studies have found that agricultural productivity is relatively lower on plots managed by women and, more importantly, that these gender gaps in agricultural productivity stem from differential access to and use of agricultural inputs among women and men (Doss, 2018; Quisumbing and Pandolfelli, 2010; Quisumbing, 1996).

While highly significant for understanding gender gaps in agriculture, the insights from these studies do not readily apply to the joint agricultural production systems that characterize most agricultural households worldwide, including in South Asia, the setting for this study. A reflection of the difference in agricultural production systems between sub-Saharan Africa and South Asia is in the share of paid and self-employed men and women workers in agriculture. In sub-Saharan Africa, 97 percent of women's agricultural labor force was self-employed (unpaid) in 2000, whereas in South Asia only 53 percent of women working in agriculture were self-employed in 2000 (FAO, IFAD, and ILO, 2010). This trend also applies to men: 93 percent men in sub-Saharan Africa were self-employed in agriculture, in contrast to 60 percent in South Asia. This pattern suggests self-employed and hired women and men are separate groups of workers in joint production systems, who perform specific tasks on the farms, face different constraints, and presumably contribute differently to agricultural productivity. This division further implies that gender-based inequities may influence hired and self-employed family women farmers differently, as compared to hired and family men farmers. In these contexts, a more relevant approach to estimating the size and consequences of gender-based gaps in agricultural productivity involves understanding differences that arise from the distinct use of family versus paid women workers, as compared to the differential use of family and paid male workers.

This study examines gender-based agricultural productivity gaps arising from the differential use of hired and family female labor, relative to the gap arising from the differential use of hired and family male labor in India, an agricultural production system in which men and women farm jointly and men primarily make farm decisions.[1] Implicitly, we study if gender-based inequities influence these labor groups differently in a way that it manifests as forgone agricultural production. We test the hypothesis that gender-based inequities influence family and hired female labor more than family and hired male labor by comparing agricultural productivity differentials between households that employ family versus hired women, relative to households that employ family versus hired men.

[1] In 98 percent of male-headed households in our study's sample, men were reported as making agricultural decisions on all the plots the household cultivated.

We interpret the resulting agricultural productivity gap as a more relevant indicator of the gender gap in productivity in joint production systems. In India, gender-based inequities influence the labor choices of rural women more than those of men (Jayachandran, 2015). The interplay between gender, wealth, and caste guides whether women are involved in farm and off-farm wage work. This may lead to differences in the skills of hired and family women workers compared to men in carrying out specialized tasks on the farm, such as transplanting, weeding, and harvesting. A farmer's skill level can be influenced by various factors, such as years of experience and the breadth of opportunities for performing specific tasks. Family women workers could have lacked such opportunities for skill acquisition due to the prevailing cultural norms dictating who they interact with and the spaces they access, disproportionately more than family male members (Quisumbing and Pandolfelli, 2010). Moreover, whereas human capital acquisition tends to make male workers more productive on the farm, more educated women are less likely to engage in farm work (Eswaran, Ramaswami, and Wadhwa, 2013). Another dimension contributing to the productivity gap between hired and family workers, in general, is the moral hazard problem of shirking. However, due to gender-based labor market frictions, this moral hazard problem could manifest differently for male and female wage workers. Gender-based inequities may imply that fewer work opportunities are available for women workers compared to men. The risk of losing wage work and not finding other commensurate work may be higher for women than men, leading them to engage in less shirking while performing agricultural wage work. Taken together, these factors suggest that a gender-based agricultural productivity gap could exist based on the household's differential use of family and hired women workers, as compared to a differential use of family and hired men farmers.

To measure the agricultural productivity gap, we analyze plot-level data obtained from households engaged in rice cultivation in the northeast Indian state of Bihar. We classify households into three groups: those relying exclusively on family female labor, those employing both family and hired women workers, and those exclusively using hired female labor.

We identify agricultural productivity differences after controlling for input use, and plot-and household-level characteristics that could influence agricultural productivity and use village-level fixed effects. To mitigate concerns regarding the non-random selection of households into these labor-use categories and potential agricultural production model misspecification, we employ machine learning estimators for measuring the agricultural productivity gap (Farrell, 2015; Chernozhukov et al., 2018; Belloni, Chernozhukov, and Hansen, 2014). Further, to infer how these average differences manifest across the entire yield distribution, we utilize a semi-parametric decomposition method (DiNardo, Fortin, and Lemieux, 1996). We construct a counterfactual density of productivity, representing the productivity that would result for households using hired female labor if they had the same observable characteristics as households using only family female labor. As a placebo test, we construct analogous household groups based on the differential use of male labor: households exclusively using family male labor and those also using hired male labor.[2]

Based on this identification strategy, we find that, on average, households using only family female labor have a lower agricultural productivity compared to those also using hired women workers. In contrast, we do not find evidence of such a gap between households using only family male labor and those also utilizing hired male labor. These results are robust to using machine learning-based variable selection and matching estimators. Across all model specifications, the magnitude of the agricultural productivity gap based on differential women's labor use is such that the value of forgone agricultural production is greater than the cost of hiring women, representing an overall welfare loss for the household. This agricultural productivity difference potentially stems from differences in opportunities to work on the farm among family and hired women. Specifically, the agricultural productivity gap is evident when comparing households where women work exclusively on their own farm to those also using hired female labor. The gap becomes much smaller when comparing households using

[2]This classification for male labor use differs slightly from that of female labor-use because more than 99 percent of our sample uses family male labor, making it difficult to construct a household group that only uses hired male labor.

only family female labor but when these women also work outside to those using hired female labor also.

Our results relate most closely to the emerging literature on the relevant aspects of gender-based gaps in agricultural productivity in joint agricultural production systems (Mahajan, 2019; Seymour, 2017). In the Indian context, Mahajan (2019) has examined the gender gap in agricultural productivity based on the sex of the farm manager, using within-village variation in farm management instead of the commonly employed within-household variation in plot management. The study reveals that women-managed farms in India have lower productivity compared to male-managed farms. Similar to our results, it presents suggestive evidence that women's lack of experience as farm managers contributes to this gender productivity gap. A key aspect of their work is that households where women manage all farm plots represent a significantly lower share of all farming families in India and are generally not the norm. These households likely differ in many ways from households where men manage all the plots.[3] Our study contributes to this literature by exploring the role of gender-based inequities among households in which men predominantly manage the farm enterprise and which represents the majority of agricultural households in India.

In Bangladesh, Seymour (2017) demonstrates that when the empowerment gap between men and women residing in the same household reduces, the technical efficiency on the farm increases. Although not directly addressing empowerment, our study shows how a skill gap between hired and family women workers may lead to an agricultural productivity gap. Skills serve as a proxy for human capital, a crucial component of women's bargaining power. The results relate to Foster and Rosenzweig (1996)'s finding that more educated farmers are able to take greater advantage of technical innovations in agriculture. We find suggestive evidence of the converse scenario: a lack of experience in performing agricultural tasks contributes to lower agricultural productivity among family women workers.

[3] In Mahajan (2019), women-managed farms constitute only 7 percent of the total agricultural households, and is based on a nationally-representative household survey.

More importantly, these findings hold significant implications for informing the policy debate on rural women's declining labor market participation in India, against the backdrop of increasing agricultural productivity (Raveendran and Kannan, 2012). From 2004-05 to 2009-10, approximately 18 million men left the agricultural sector but found employment in other sectors, such as construction. In contrast, about 36 million women left agriculture but did not engage in non-farm employment. Most of these women who exited agriculture were self-employed women working on their own farms (Rangarajan, Kaul, and Seema, 2011). This study shows that family women's involvement on the farm is linked to lower agricultural productivity compared to hired women workers. On one hand, the exit of self-employed women from the agricultural sector suggests a potential improvement in agricultural productivity. On the other hand, if women's skills are perceived as more suitable for unpaid household work than unpaid farm work, it not only threatens women's bargaining power in rural labor markets but also perpetuates restrictive gender and social norms (Eswaran, Ramaswami, and Wadhwa, 2013). Considering these trade offs, it would be beneficial to direct investments towards expanding women's opportunities in the rural labor market, both on- and off-farm. Moreover, holistic programs aimed at removing gender-based barriers restricting women's labor choices would prove beneficial.

## 2 Conceptual Framework

This section illustrates a framework of how agricultural productivity gaps may emerge in joint agricultural production systems. Similar to studies that compare productivity differences between women- and men-managed plots within a household, we postulate that agricultural output, $q$, is produced using inputs and is represented as follows.

$$q = G(\cdot) \tag{1}$$

Here $G(\cdot)$ is a concave and increasing production function in inputs. Most studies examining the gender gap in agricultural productivity test the equality of output produced by different plot managers within a household, conditional on controlling for plot-specific factors that could influence productivity. That is, they test if $G_m$ and $G_w$ are equal for plots managed by men ($m$) and women ($w$) within a household for similar-quality plots (for example, Udry (1996); O'Sullivan et al. (2014). However, in joint production systems, a clear gendered demarcation in plot management does not exist, and, often, one single farm manager makes the input allocation decision across all plots. That is, in our context, we consider a single production function, $G(\cdot)$, such that production is managed by the farm manager, primarily men in male-headed households.[4]

In joint agricultural production systems with a single agricultural production function, one could measure women's and men's labor productivity, but any such differences may not be meaningful because women's and men's labor are often complements instead of substitutes in production (Doss, 2018). More importantly, if such analysis shows, for example, that women's marginal product is higher in transplanting relative to men and men's marginal product is higher in operating machinery relative to women, then a strengthening of such roles may actually worsen gender inequities. Instead, gender-based inequities in such contexts may manifest as differential productivity between hired and family women's labor vis-à-vis men. Specifically, in the Indian context, both caste and gender norms could influence women's labor choices, mobility, and access to technical know-how more than men. As such, this could create a distinction between the capabilities of hired and family women, that is significantly more pronounced than hired and family men, which could, ultimately, translate to generating agricultural productivity gaps that arise from gender inequities, but are not gendered in the traditional sense of comparing differences between women's and men's contributions to agricultural production. From both a conceptual and policy perspective,

[4]This set-up closely mimics the agricultural production system broadly prevalent in India. In our study's context, men managed all plots in the household in 97 percent of all our study's sample and in 98 percent of male-headed households.

understanding and addressing such gaps in joint agricultural production systems can be more meaningful, both within individual households and across them. In the framework and the ensuing analysis, we focus on understanding how such agricultural productivity gaps may appear.

In our set-up, we assume agricultural production uses three inputs: (i) women's labor, $E$; (ii) an aggregate measure of all other inputs, such as seeds, fertilizer, and men's labor, represented by $M$; and (iii) land, which is fixed and denoted by $D$. For simplicity, we only consider women's labor as a separate input and subsume men's labor in the aggregate measure of all other inputs. The output, $q$, is written as follows.

$$q = G(M, E, D) \tag{2}$$

A key dimension of the framework is the notion of "effective" women's labor (Eswaran and Kotwal, 1985). We assume households can use women's family labor, represented by $L_f$, or hire women's hired labor, denoted by $L_h$. However, the actual labor used on the farm, that is $E_f$ and $E_h$, is different from $L_f$ and $L_h$. Effective hired women's labor is defined as:

$$E_h = \gamma_h(\sigma_h, s_h, c_v)L_h \tag{3}$$

Here, $\gamma_h(\cdot)$ is a concave efficiency parameter, increasing in three factors. The first factor, $\sigma_h$, captures human capital or experience of women workers, an important aspect of worker productivity (Sahn and Alderman, 1988). In our context, for women workers who are likely not making many managerial decisions, human capital is embodied not so much in the level of education but in experience performing specific rice cultivation tasks and opportunities for learning the appropriate skills needed to perform tasks. The more experience women workers have in performing specific cultivation tasks, such as transplanting, weeding, and harvesting, the more productive they would be on the farm. The second factor that influences productivity is time allocated to supervision by the farm manager, represented as $s_h$ (Eswaran and Kotwal, 1985). Because the moral hazard problem of shirking is an inherent problem in hiring workers, we assume all hired workers are supervised, and that worker productivity is increasing in supervision (although with diminishing returns to additional units of supervision). The third factor that influences hired women workers' quality of labor is the cost of finding alternative work for women, in case they are dismissed due to poor performance or other related factors. We represent this cost as $c_v$, which we assume to be constant at the local-economy, village-level (denoted by subscript $v$), in which women reside and seek work.[5] Higher the cost of finding work, lower would be the likelihood of women workers shirking, and, consequently, higher would be their productivity.

Similarly, effective family women's labor is represented as follows.

$$E_f = \gamma_f(\sigma_f)L_f \tag{4}$$

As shown in Equation 4, family women's efficiency only depends on their experience, $\sigma_f$. Because there is no moral hazard problem of shirking associated with household family members supplying labor on their own farm, they do not need to be supervised. Moreover, we also assume that family women only supply unpaid labor on their farm and do not work on other farms as hired workers.

Our next set of assumptions pertain to prices. We assume that the opportunity wage of family women workers is represented by $u$, of family men supervising hired workers by $\delta$, and and of hired women workers by $w$. All opportunity wages are assumed to be exogenous. The price of the aggregate measure of all other inputs is $p$. The price of output is unity. All workers and male supervisors have one unit of time, which they allocate between agricultural production and alternative activities. We also assume labor markets of hired women workers

[5] Implicitly, we assume that women workers do not migrate alone to other places in search of work opportunities and only seek wage work in the vicinity of their homes. This assumption about the extent of women's labor market is valid in the rural Indian context (Luke and Munshi, 2011)

are non-missing and competitive.[6]

We now turn to examining the optimization problem facing the household and consider two scenarios based on the differential use of women's labor ($E$): household only uses either hired women workers or only family women's labor.

**Case 1: Use only Hired Women's Labor.** The optimization problem is illustrated as follows.

$$\max_{s_h, M, L_h} \quad [G(M, \gamma_h(\sigma_h, s_h, c_v)L_h, D) - pM - wL_h] + (1 - s_h)\delta$$

The first order condition with respect to hiring women workers yields the condition that the "effective" marginal product of hired women workers equals the marginal cost of hiring them. That is,

$$\gamma_h(s_h, c_v)\frac{\partial F(\cdot)}{\partial L_h} = w \tag{5}$$

**Case 2: Use only Family Women's Labor.** When the household only uses family women's labor, it pays the opportunity cost of their time and purchases other inputs.

$$\max_{s_f, M, L_f} \quad [G(M, \gamma_f(\sigma_f)L_f, D) - pM - uL_f] + \delta$$

The first order condition for using only family women's labor also yields a similar condition and is represented as:

$$\gamma_f(\sigma_f)\frac{\partial F(\cdot)}{\partial L_f} = u \tag{6}$$

Using Equations 5 and 6 as the foundation, we now consider two types of asymmetries, rooted in gender inequities and norms, that could results in an agricultural productivity gap based on differential women's labor-use and not based on men's labor.

**1. Gender- and caste-based inequities in labor supply.** In the Indian context, the

[6]This assumption is valid in our context because in all villages in our sample, at least a few households hired women workers, with not a significant variation in wage rates.

intersection of wealth, caste, and gender may govern the labor supply decisions of women (Luke and Munshi, 2011; Paris et al., 2000; Chen, 1989). Women from land-poor, low-caste households are more likely to work as hired workers as compared to women from upper-caste wealthier households. Even when women can work outside, their mobility might be constrained, limiting the rural farm- and off-farm jobs they can perform. These restrictions on their work choices and mobility may influence both the level of experience in farming ($\sigma$) and the cost of finding work in their local communities ($c_v$). Specifically, because hired women workers presumably have had greater opportunities to perform specific specialized tasks, such as transplanting paddies, weeding, and harvesting, as compared to family women, $\sigma_h$ might be higher than $\sigma_f$. Moreover, hired women workers may also have more opportunities to learn from one another. In contrast, family women may not have been exposed to many learning opportunities, from other women in their family or their friends. Such a skill-based gap may only emerge in hired and family women's labor and not in hired and family men's labor because family men may are also farm managers, have greater mobility, and may have greater access to information and technical know-how. All else being equal, these asymmetries may imply that the effective marginal product of hired women's labor might be higher than family women's labor, which may lead to agricultural productivity gaps among households that only use hired women's labor as compared to family women's labor.

**2. Gender-based asymmetry in input allocation decisions.** As our data suggests and as is prevalent in most patriarchal joint production system, men primarily make input allocation decisions. Whereas supervision of hired workers, $s_h$, may allow them to have an accurate belief about $\gamma_h(\cdot)$, a lack of such supervision of family women may imply they have inaccurate beliefs about family women's productivity ($\gamma_f(\cdot)$). This is especially problematic if they perceive that family women's productivity is equivalent to hired women. This lack of supervision and a misperception about $\gamma_f(\cdot)$ may also result in agricultural productivity gaps based on the differential use of women's labor, which may not hold for men because of the observability of family and hired men's productivity.

In the empirical analysis that follows, a key identification challenge is that we only observe the household's actual agricultural output and the amount of family and hired women workers it uses. We do not observe the changes in output had the household employed a different mix of family and hired women workers. To circumvent this issue, as a first step, we classify the households in our sample based on using family women workers only, both family and hired women workers, or hired women's labor only. We then employ both fixed effects and machine-learning based matching methods to measure agricultural productivity differences for households that appear to be alike on observable characteristics but differ on their use of family and hired women workers. Whereas our empirical data does not allow us to understand the role of the second asymmetry, we indirectly examine if the agricultural productivity gap is associated with a skill-gap between hired and women workers and women's local labor market.

# 3 Context and Data

The data for the study were collected as part of a broader research program on understanding gendered division of labor in rice cultivation in the north Indian state, Bihar. Rice and wheat are two key staple crops cultivated in India, and women farmers contribute a greater share of their labor to rice cultivation, as compared to wheat (Boserup, 1970; Rosenzweig and Schultz, 1982; IRRI, 1983). In addition to family women's labor-use, women laborers are also hired for transplanting, weeding, and harvesting rice, as shown in Figure 1. Rice cultivation, thus, provides an appropriate context for understanding agricultural productivity differences arising from using family and hired women's labor.

In 2015, a month after rice harvest, we collected plot-level, gender-disaggregated labor and capital use data on rice cultivation from randomly-selected agricultural households that had at least one adult male and female co-decisionmaker. We also collected data on the household's socio-demographic information, along with information on the adult male and female co-decisionmaker's labor market participation. Although data from 965 households was collected for the broader research program, this study restricts the sample to households reporting a non-zero value of rice output and households where output productivity is not a significant outlier, which could stem from errors in data collection or reporting (Wollburg, Tiberti, and Zezza, 2021).[7] Our final sample consists of 681 households drawn randomly from 27 villages in 13 districts.

As discussed in Section 2, a key distinction we make pertains to the household's use of women's hired and family labor. As the basis of our analysis, we classify sample households into three groups based on the type of women's labor they use in rice cultivation. The first group comprises households that only use family women's labor for cultivating rice (referred to as "only family women"). The second group of households use both family and hired women's labor during rice cultivation (referred to as "both"). The third group of households rely only on using hired women's labor or no women's labor for cultivating rice (referred to as "only hired women"). As Table 1 shows, 27 percent of sample households belong to the "only family women", 33 percent in the "both", and 39 percent in the "only hired women" group. Throughout the remainder of the paper, we present analysis based on the classification of women's labor use and test the hypothesis that "both" or "only hired women" households are associated with higher rice output productivity, as compared to "only family women" households. Moreover, in order to establish that these differences exist only based on the differential use of hired and family labor among women and not men, we also implement a placebo test based on a similar household classification for men. We classify households into those employing only family men's labor and those that also use hired men workers. This classification differs from that of women because only 1.5 percent of sample households do not use any family men's labor and rely exclusively on hired men's labor, making it difficult

[7] We restrict the sample to households whose rice output per acre is below or equal to the 95th percentile of the productivity distribution. We also remove outliers (restricting values to below the 99th percentile) based on plot size and key inputs, such as total men's and women's labor used per acre, hours spent irrigating the field per acre, and the amount of fertilizer (urea) used per acre.

to construct this household group for men.

## 3.1 Descriptive Statistics

Table 1 provides a snapshot of the socio-demographic and plot-level characteristics of the sample households. Column (1) shows the summary statistics for households that use "only family" women for cultivating rice, and columns (2) and (3) show the average deviation from this base group for households using "both" hired and family women's labor and those using "only hired" women's labor. As the table suggests, these groups, on average, vary in few key household characteristics. Although most households in our sample are headed by men (about 98 percent), households that use only hired women's labor are almost entirely headed by men. The household head also has about 5 additional years of average education in households that use only hired women's labor. In the Indian context, the intersection of wealth, caste, and gender guides the work rural women perform, and our data also supports this perspective descriptively (Eswaran, Ramaswami, and Wadhwa, 2013; Jayachandran, 2015). Among households that use only family women, only 9 percent households are upper caste. This share also is similar for households that use both family and hired women's labor. In contrast, more than 55 percent households are upper caste in the "only hired" group. The two household groups that hire women also have a higher wealth index, on average, as compared to households that rely on only family women's labor.[8] Among all the household groups, there are no statistically significant differences in the household head's age (about 48 years), the household size (about 6 members), and religion (primarily Hindu). Despite differences in wealth index, the households also do not exhibit statistically significant differences in their ability to obtain loans, which is presumably an important determinant of the inputs households use.

[8]The wealth index is constructed using factor analysis, and is based on ownership of the following assets: cellphones, motorcycle, television units, cable television, diesel pump, rotavator, knapsack, and tractor; and the size of land owned (in acres). Further, it also includes household's reported expenditures on transport, education, and festival donations.

Next, we examine the unconditional differences in crop cultivation patterns among these households. Households that exclusively utilize hired women's labor farm on larger plots, which they have cultivated for a longer duration, compared to households that solely rely on family women's labor. They are also more likely to own the plots they cultivate as compared to households that use only family women's labor. These plot characteristics, such as size, duration of cultivation, and tenancy are not statistically different between households that use only family women and households that use both family and hired women. There are no differences in the share of plots being upland among the three household groups.

Turning to input use, on average, households that use only family women's labor use about 31 days of women's labor and 57 days of men's labor per acre. Households that employ both family and hired women's labor use about 6 days more of women's labor per acre, but this difference is not statistically significant. In contrast, households where only hired women work use about 12 labor days less for cultivating an acre of land. Although households that use any hired labor (that is, the "both" and "only hired" women groups), use about 10 days less of male labor per acre, this difference is not statistically significant. Fertilizer use – both in terms of the quantity of urea used and the percentage of households using any non-urea fertilizer – does not vary significantly across groups. However, households that hire any women irrigate their plots longer as compared to households that use only family women's labor. These households are also likely to own or hire a tractor and a diesel or an electric pump for irrigation. Overall, the use of pesticides, herbicides, and fungicides in our sample is low and there are no statistically significant differences in their use among the three household groups. These differences imply that households that are using hired women's labor are also presumably using different production technologies in crop cultivation, which could influence agricultural productivity. Therefore, we control for these key differences in plot characteristics and input use in our empirical analysis.

## 3.2 Agricultural Productivity and Women's Labor

Table 1 shows the unconditional agricultural productivity, measured by kilograms of rice produced per acre, is lowest for households that use only family women's labor. Compared to households using only family women's labor, agricultural output is higher by 231 kilograms per acre in households using both types of women's labor and by 248 kilograms per acre in households using only hired women's labor.[9] In other words, agricultural productivity in households that employ any hired women's labor is at least 21 percent higher than the average agricultural productivity in households that rely on only family women's labor, with the differences being statistically significant.

To further examine whether these yield differentials hold for different plot sizes, Figure 2 illustrates the non-parametric relationship between yield (measured as a deviation from the village average yield) and plot area. As the figure suggests, the yield differential between households that use only family women's and those that use any hired women's labor remains consistent for almost all plot sizes. The gap in yield becomes smaller as plot area increases. The figure also suggests that similar to other studies in the literature, plot size and agricultural productivity are inversely related in our study's setting as well (Bindlish et al., 1993; Udry, 1996). In the following sections, we examine the conditional agricultural productivity differences among these household groups.

# 4 Empirical Strategy

## 4.1 Fixed Effect Estimation

To measure the agricultural productivity differences arising from the type of women's labor used, we estimate an agricultural production function based on the inputs used, as well as

[9] According to official statistics, rice yield in India in 2015-16 was about 1322 kilograms per acre, which is similar to the rice yield reported among households using only hired women's labor (United States Department of Agriculture, 2023).

the plot- and household-level characteristics. We include the three household groups as a categorical variable, with households using family women's labor as the omitted group. We implement this to identify agricultural productivity differences arising from the differential use of women's labor after controlling for differences in inputs and plot and household-level characteristics. Equation 7 shows the main specification.

$$Y_{ihv} = \beta_0 + \beta_1 LaborType_{hv} + \beta_2 P_{ihv} + \beta_3 I_{ihv} + \beta_4 H_{hv} + V_v + \epsilon_{ihv} \tag{7}$$

The outcome, $Y_{ihv}$, represents agricultural output productivity on plot $i$ cultivated by household $h$ in village $v$, measured as the logarithm of the quantity of rice produced (in kilograms) per acre. The main coefficient of interest is associated with $LaborType_{hv}$, a categorical variable that captures whether the household uses hired and family women's labor or only relies on hired women's labor. The omitted group includes households that use only family women's labor. We employ a fixed effects approach at the village level, as captured by $V_v$. This strategy allows us to control for unobservable characteristics across local geographies, such as agro-climatic and market access variation, that could influence input use and crop productivity at the village level.

In our estimation, we include a wide range of plot characteristics, inputs used, and household-level variables that could influence agricultural production. Specifically, $P_{ihv}$ captures plot characteristics, such as size, soil type, ownership status (share cropped or rented in), years of cultivation, and slope. $I_{ihv}$ includes agricultural inputs used on each plot: hours of female and male labor used per acre, hours of irrigation per acre, and the amount of fertilizer used per acre. We include the logarithmic transformation of these variables. Further, $H_{hv}$ captures household characteristics that could influence agricultural productivity, such as age and education of the household head; whether the household head is male; whether the household is upper caste and Muslim; household size; whether the household is eligible to receive a loan; whether the household uses any pesticide, herbicide, or fungicide, and whether the household owns or rents a tractor and a diesel or an electric pump for irrigation.

## 4.2 Machine Learning Estimation

A key assumption we make in the empirical strategy described above is that, conditional on including the relevant plot-level ($P_{ihv}$), input-use ($I_{ihv}$), household-level ($H_{hv}$), and village dummy ($V_v$) variables, the household's use of the type of labor ($LaborType_{hv}$) is exogenous (the *unconfoundedness* assumption) and unrelated to the error term (the *conditional independence* assumption). Whereas the use of village-level fixed effects estimator and extensive control variables helps address some of these identification concerns, these assumptions are still non-testable. Therefore, we employ two machine learning-based approaches to account for potential model misspecification and non-random selection of households into these labor-use categories.

For simplicity, we denote the plot-, input-, and household-level variables as $g(x)$. That is, Equation 7 is represented as:

$$Y_{ihv} = \beta_0 + \beta_1 LaborType_{hv} + g(x) + V_v + \epsilon_{ihv} \tag{8}$$

As the first approach to obtaining a better estimate of $g(x)$, we employ a partial linear model that uses the post-double selection least absolute shrinkage and selection operator (PDS LASSO), as proposed by Belloni, Chernozhukov, and Hansen (2014). The PDS LASSO estimator allows for appropriately selecting control variables to improve precision and accuracy of the estimates obtained for $LaborType_{hv}$ and helps improve the robustness of the estimates. This estimator is especially relevant in our setting because plot-, input-, and household-level variables could interact in numerous ways in an agricultural production function and contribute to agricultural productivity differences across households. This approach helps in selecting appropriate controls from an exhaustive list of controls, including squaring all variables and interacting them with one another. Appendix A1 shows the list of all variables that are standardized and included for selection using the PDS LASSO estimator.

The second machine learning tool we employ uses a heterogeneous effects model estimator, the augmented inverse probability weighting LASSO (referred to as AIPW LASSO) (Chernozhukov et al., 2018; Farrell, 2015; Belloni, Chernozhukov, and Hansen, 2014; Cameron, 2022). This approach estimates a separate model for agricultural productivity on $g(x)$ for the two household categories, households using family women's labor and those using hired women's labor) using the LASSO procedure. Then, to account for potential selection of households into each of these labor-use categories, it uses the LASSO method to select variables that predict the $LaborType$ variable. Using these variables, it estimates the conditional probability or the propensity score indicating the likelihood of a household belonging to one of the labor-use categories. Using these weights that account for the probability of selection of a household into a labor-use category, the estimator computes the weighted average of predicted agricultural productivity of households using hired women's labor and those using family women's labor. The difference of the two weighted averages provides an estimate of the agricultural productivity gap. The key value addition of this estimator, as compared to PDS LASSO, is that it is doubly-robust to model misspecification and is consistent either if the outcome model or the propensity score model is misspecified, thereby providing additional flexibility. This estimator further enhances the plausibility of the *unconfoundedness* and the *conditional independence* assumptions.

## 4.3 Semi-Parametric Decomposition

The fixed effect and machine learning estimation methods allow us to measure the average agricultural productivity gap based on the household's use of family and hired women's labor. To understand how this gap is distributed along the agricultural productivity distribution, we estimate a semi-parametric decomposition proposed by DiNardo, Fortin, and Lemieux (1996). This decomposition method is analogous to the linear Kitagawa-Blinder-Oaxaca decomposition technique (Kitagawa, 1955; Blinder, 1973; Oaxaca, 1973). In the semi-parametric decomposition, we construct a counterfactual agricultural productivity for households using hired women's labor but conditional on having the same characteristics as households using only family women's labor. We compare this counterfactual density with the actual distribution of productivity for households using only family women's labor. Put simply, this decomposition method allows us to understand the agricultural productivity that households using family women's labor would achieve if they used hired women's labor instead, while holding all other observable characteristics constant.

Using similar notation as Mahajan (2019), the actual agricultural productivity density functions of households that use family and hired women's labor are represented as:

$$f(y|L=F) = \int g(y|x, L=F)h(x|L=F)dx \tag{9}$$

$$f(y|L=H) = \int g(y|x, L=H)h(x|L=H)dx \tag{10}$$

In Equations 9 and 10, $L$ represents the type of women's labor used by the household, family ($F$) or hired ($H$). Functions $f(\cdot)$ represents the unconditional productivity density, $g(\cdot)$ the productivity density conditional on a vector of observable household characteristics ($x$), and $h(\cdot)$ the density of observable characteristics conditional on the household labor type. The counterfactual that we aim to construct is the density that would prevail for households if they used hired women's labor (that is, $L = H$) but had the same characteristics as households using family women's labor (that is, vector $x$ conditional on $L = F$). This counterfactual density is shown as:

$$f_{counter}(y) = \int g(y|x, L=H)h(x|L=F)dx \tag{11}$$

We construct the counterfactual density shown in Equation 11 and compare it to the actual density of households using family women's labor shown in Equation 9 to understand how productivity would be distributed if households used hired women's labor instead of using family women's labor.

Put together, these estimation strategies allow us to increase the plausibility of the conditional independence and unconfoundedness assumptions related to the consistent estimation of the coefficient associated with the *LaborType* variable. The machine learning estimators help ensure that we include permutations of input-use variables and farmer-specific factors that may lead to agricultural productivity heterogeneity, typical in crop cultivation. This approach also helps minimize the bias in our estimates arising from us selecting variables to include in the regression analysis. Moreover, we also estimate the productivity gap in the entire agricultural productivity distribution using a semi-parametric decomposition, which helps us to understand the part of the distribution in which the gap is the largest. More importantly, as a placebo test, we estimate the productivity gap based on the household's differential use of male labor. We test the difference in agricultural productivity between households using only family men's labor and those also using hired men's labor. Further, Section 7 checks the sensitivity of our results to omitted variable bias stemming from using observational data and non-random assignment of households into each labor type (Oster, 2019; Altonji, Elder, and Taber, 2005).

# 5 Results

## 5.1 Main Results

Table 2 shows the results associated with estimating Equations 7 and 8. Households using only family women's labor represent the omitted group. Panel A shows the estimated coefficients associated with households using only hired women's labor and those using both, and Panel B presents these coefficients as percentage changes. Columns (1), (2), and (3) show the regression coefficients obtained using the village-level fixed effect estimator and with the inclusion of plot-, input-, and household-specific control variables that could influence agricultural productivity.[10] The results suggest that agricultural productivity, as measured by the quantity of rice produced per acre, in households that use only hired women's labor is higher by at least 37 percent as compared to households that use only family women's labor, and is statistically significant. Households that use both hired and family women's labor have a higher agricultural productivity by at least 18 percent as compared to households that only rely on family women's labor, although the regression coefficient associated with households using "both" family and hired women's labor is statistically insignificant. The magnitude of the estimates associated with the "both" and "only hired women" groups does not change drastically across the different regression specifications. Moreover, the difference in point estimate associated with the "both" and "only hired women" household groups is also statistically significantly and both estimates are jointly significant.[11]

Turning to the machine learning models, columns (4), (5), and (6) show the coefficients associated with using the PDS LASSO estimator for selecting the control variables. The magnitude of the coefficient associated with households that only use hired women's labor is lower compared to that obtained in the fixed effect model. This suggests that there may be other observable characteristics accounting for variation in agricultural productivity that were not considered in the fixed effect model. Agricultural productivity in households using only hired women's labor is at least 28 percent higher compared to households using only family women's labor, based on the PDS LASSO estimator. However, the estimated coefficients associated with households using both hired and family women's labor are statistically significant when we use the PDS LASSO estimator, and the magnitude is similar to that obtained in the fixed effect estimation (agricultural productivity is higher by approximately 16 percent among this group).

When we account for the household's non-random selection into these labor-use categories

[10] The complete results obtained from the fixed effect estimation are shown in Appendix A2.

[11] As Appendix A2 shows, across the three fixed effect regression specifications, the wealth index of the household is negatively associated with agricultural productivity, which could reflect the inverse land size and agricultural productivity relationship. The coefficient associated with female labor use is also statistically significant and is associated with increasing agricultural productivity.

using the AIPW LASSO estimator, the magnitude of the coefficients associated with using only hired labor or using both family and hired labor increase.[12] Households using only hired women's labor have an agricultural productivity that is about 46 percent higher compared to households using only family women's labor. Furthermore, agricultural productivity in households using both hired and family women's labor is higher by as much as 55 percent compared to households using only family women's labor. In other words, when we apply inverse probability weights to observations and obtain labor-use groups that are balanced conditional on observable characteristics selected using the LASSO procedure, the agricultural gap increases in size between households using family and any hired women's labor. According to the AIPW LASSO estimator, the agricultural productivity gap is highest between households using both hired and family women's labor and those using only family women's labor. Presumably, hired women in the "both" group are performing the majority of specialized tasks, and family women in this group could be assuming more supervisory roles.

Recall, the average rice productivity of households using only family women's labor is 1085 kilograms per acre, and the unconditional agricultural productivity gap between households using only hired and only family women's labor is 23 percent. The regression results obtained for all the three models suggest that the conditional agricultural gap is at least as high as the unconditional gap and could be as high as 46 percent when we control for non-random selection and omitted factors that could influence agricultural productivity. Similarly, the unconditional agricultural gap between households using both hired and family's labor is 21 percent, and the regression estimates suggest that the conditional gap could be slightly lower about 16 percent percent) but could be as high as 55 percent when we account for selection and confoundedness in our estimates.

More importantly, the size of the agricultural productivity gap is economically significant:

[12] Given that the AIPW LASSO estimator creates two groups that are balanced conditional on observable characteristics, there is no regression specification that includes both the "only hired" and "both" groups together.

The value of output per acre lost more than offsets the cost of hiring women workers, on average. As shown in Table 1, the mean and mode output price of rice was Indian Rupees (INR) 46 and 10 per kilogram, respectively, at the time the data were collected. Using the conservative mode price, this implies that the value of output lost for households using only family women's labor is at least INR 3079 per acre compared to households using only hired women's labor.[13] The mode women's daily wage rate in the sample was INR 125 and INR 150 for transplanting and weeding, respectively. If we use a wage rate of INR 150, this implies that households could hire at least 21 labor days of hired women workers, assuming all other costs are equal, which is higher than the average labor days of hired women workers used in rice cultivation per acre in our sample for households using only women's labor. That is, in the most conservative scenario, the magnitude of the gap is such that the household could hire at least as many hired women workers as the cost of foregone production.

Table 3 shows the results from the placebo test obtained from using the labor-use categories for men.[14] Because there are very few households that do not use family men's labor, we categorize the households into those using "only family men's" labor and those also using "hired men's" labor. We reject the existence of significant agricultural productivity gaps based on using hired men's labor: The regression coefficients associated with also using hired men's labor are low in magnitude and statistically insignificant using the fixed effect, PDS LASSO, and AIPW LASSO regression models. The magnitude is the highest (agricultural productivity gap is 11 percent higher) when we use the PDS LASSO model, which accounts for potential non-random selection of households into each of the labor-use categories and also selects control variables using the LASSO method. The size of the gap is lower than the lowest agricultural productivity gap obtained for women's labor-use.

Appendices A4 and A5 show the robustness of our main results for samples restricted to male

[13] This refers to a 28 percent agricultural productivity gap, which is the lowest point estimate associated with households using only hired women's labor obtained across the fixed effect, PDS LASSO, and AIPW LASSO regression models.

[14] Appendix A3 shows the full table of regression estimates obtained using fixed effects model based on the household's use of hired and family men.

headed households and households belonging to the lower caste, respectively.[15] In households headed by men, family women's roles could be different as compared to households headed by women where women may be performing more managerial roles. Because we are interested in examining differences based on labor used, we restrict the sample to male headed households. The results suggest that the point estimates are similar to those shown in Table 2. Similarly, although a high share of households are classified as belonging to the lower caste in our sample, it could be the case that this productivity gap is being driven by family women belonging to the upper caste. The results suggest that although the magnitude of the gap is slightly lower, we still find statistically and economically significant evidence of an agricultural productivity gap between households using family and hired women's labor in the lower-caste households. Appendix A6 also shows the agricultural productivity gap using a household fixed effect model, which exploits the within household variation in the use of different kinds of women's labor. Although household fixed effects are more robust than using a fixed effect model at the village level because they control for any omitted household-level characteristics that could influence agricultural productivity, the actual number of households in our sample that have such variation is low. We report these results as a check for robustness. Using a household fixed effect model, the size of the agricultural productivity gap associated with not using family women's labor on a given plot is 21 percent and statistically significant.

## 5.2 Semi-Parametric Decomposition

Figure 3 shows the results obtained from the semi-parametric decomposition in which we construct the counterfactual agricultural productivity distribution that would prevail if households used hired women's labor instead of family women's labor. Part (a) in the figure shows the kernel density of agricultural productivity of households that use only family women's labor. The counterfactual density represents the density that would prevail for households

[15]Households that reported their caste as scheduled caste, scheduled tribe, or other backward caste were classified as belonging to the lower caste.

using only hired women's labor if they had the same observable characteristics as households using only family women's labor.[16] As the figure suggests, if households that use hired women's labor had the same characteristics as households that use family labor, the resulting productivity would be higher. The difference would be most prominent in the low-middle part of the productivity distribution for households using only family women's labor.

Part (b) constructs the counterfactual density for households that use both hired and family women's labor, under the scenario that these households have the same characteristics as those in the "only family women" group. Although the difference in the counterfactual productivity and the actual agricultural productivity distribution of households using only family women's labor is less stark as compared to the counterfactual density depicted in part (a), this household group would also have a slightly higher agricultural productivity. Appendix A7 shows the results from the semi-parametric decomposition constructed for men as a placebo test. The counterfactual density is based on the agricultural productivity of households using any hired men's labor but under the scenario that they have the same characteristics as households using only family men's labor. The figure suggests that the gap between the actual and the counterfactual density is less stark, and very similar to the counterfactual distribution obtained for households that use both family and hired women's labor. The fact that households relying exclusively on hired women's labor show a counterfactual productivity distribution that is notably skewed to the right, in comparison to the counterfactual distribution resulting from the use of hired men's labor, reinforces our findings on the agricultural productivity gap arising from differential women's labor use, and not from differential men's labor use.

[16] We construct the counterfactual density distributions conditional on the following observable characteristics: wealth index, household size, age of household head, plot size, female labor used per acre, male labor used per acre, hours of irrigation per acre, and quantity of urea applied per acre.

# 6 Potential Mechanisms

The results show that households using only family women's labor have a lower agricultural productivity compared to those using hired women's labor in cultivating rice, after accounting for variation in plot-, input-, and household-level characteristics that could contribute to agricultural productivity across these households. As discussed in Section 2, a skill and experience gap between family and hired women workers could contribute to this agricultural productivity differential. As depicted in Figure 1, family and hired women predominantly supply labor to transplanting, harvesting, and weeding during rice cultivation. All three tasks are labor-intensive and critical for rice productivity, and workers' skill and experience in performing these tasks could directly influence crop yield (Saradamoni, 1987). Women who work as hired workers could be more knowledgeable about best practices associated with transplanting, weeding, and harvesting, compared to women who only work on their own farm during rice cultivation. Hired women workers could also have had more opportunities to learn from other workers. Although we do not have data on the skill-level of hired and family women workers in performing each of these tasks, we examine this mechanism indirectly. A subset of the female co-heads in our sample who supply labor on their farms also reported engaging in paid work others' farms and off-farm.[17] We estimate the agricultural productivity gap between households that use hired women's labor and family women's labor, separately for these sub-samples of households: households in which the female co-head works exclusively on her own farm and households in which the female co-head also reports working outside.

Table 4 presents the results for these sub-samples. Columns (1) and (2) show the results from

[17] We do not have data on whether the female co-head works on others' farms during rice cultivation, hence we use the measure that captures female co-head's any paid employment.

the fixed effect model, and columns (3) and (4) from using PDS LASSO selection procedure.[18] Using the fixed effect model, we find that there is no statistically significant difference in agricultural productivity between households using only hired women workers and households using only family women and in which the female co-head also works outside. In contrast, the agricultural productivity gap is larger between households using hired women's labor and households that use only family women's labor and in which the female co-head works exclusively on her own farm. As compared to the magnitude of the gap for the whole sample shown in column (1) of Table 2, the gap is significantly higher (the size is 54 percent). When we examine heterogeneity in the size of the gap based on the PDS LASSO procedure, we find that the size of the gap is again higher for the sub-sample of households in which the female co-head does not work outside.[19] As shown in columns (2) and (4), the magnitude of the gap is closest in magnitude to estimates obtained from the AIPW LASSO full sample model (columns (7) and (8) in Table 2), which is the model that aims to account for non-random selection of households as well.

In addition to differences in skills among hired and family women, there could be other mechanisms that contribute to the agricultural productivity gap. As discussed in the conceptual framework, hired women workers could be less likely to shirk (and alternatively put in effort that is at par with family women workers) if the cost of finding alternative work is high. Given effort is difficult to measure in our context, we test this channel indirectly. We examine agricultural productivity differences based on the variation in the share of households using only family women's labor in each village. Villages in which the share of sample households using only family women's labor is low could potentially have greater availability of hired women workers to work on the farm. This aspect could imply that fewer hired women

[18]We do not implement the PDS LASSO estimation procedure separately for each of the sub-samples as that would estimate a separate model for each. Instead, for comparability and to examine the difference in magnitude of the gap, we use the variables selected by PDS LASSO estimation for the full sample to implement the sub-sample heterogeneity analysis.

[19]Note that the size of gap associated with households in which the female co-head works outside is also large (about 31 percent) but it is smaller than that obtained for those households in which the female co-head does not work outside.

workers would shirk if the risk of being replaced by others is high. If the lack of shirking by hired women workers is contributing to the gap, we would expect to find that the size of the agricultural gap is higher in villages where the share of households that employ hired workers is higher. Table 5 reports the agricultural productivity gap results for the sub-sample of households varying based on the share of households using only family women's labor in a village. In each column, we cumulatively add the villages where the share of family women is lower than 20, 40, 60, 80, and 100 percent. As the regression results suggest, the agricultural productivity gap does not change much across the different sub-samples, weakly suggesting that, perhaps, differences in shirking is likely not a potential channel. Since this is an indirect check, it is difficult to fully ascertain the contribution of this mechanism to the agricultural productivity gap.

We rely on the literature to understand the potential factors contributing to the agricultural productivity gap that could apply in our context. Mahajan (2019) finds that the agricultural productivity gap based on the gender of the farm manager is driven by women's lack of experience and knowledge in farming in India. In their study's context, a yield differential exists between men and women farm managers who cultivate wheat as compared to rice. Mahajan (2019) postulates that because women play a greater role performing rice cultivation tasks than in wheat cultivation, they could have more knowledge and experience managing rice than wheat, which could contribute to the agricultural productivity gap. The analysis in this paper additionally indicates that within households engaged in rice cultivation, differences in agricultural productivity arise depending on whether the labor is provided by hired or family women. This difference may be attributed to disparities in skill levels between family women exclusively contributing labor on their own farm and lacking experience and hired women workers, who potentially have more experience. Together, these studies suggest that in joint agricultural production systems where men are the primary decision-makers, women's lack of experience or technical skills could be a key contributor to agricultural productivity gaps. These findings are different but complementary to the existing literature on agricultural productivity gaps in sub-Saharan agricultural systems that shows that lack of access to critical inputs contributes to gender-based differences in crop yield (O'Sullivan et al., 2014; Quisumbing and Pandolfelli, 2010).

# 7 Sensitivity

As noted in Section 4, obtaining a consistent estimate of the agricultural productivity gap depends on satisfying the *unconfoundedness* and *conditional independence* assumptions. That is, conditional on the inclusion of relevant plot-, household-, and village-level control variables, the household's use of the type of women's labor, is exogenous. Although we use both village-level fixed effect and machine-learning based approaches to mitigate concerns related to the identification of the estimated coefficients, in reality, there could be both unobserved and omitted factors that guide whether households employ family or hired women's labor. For example, how household demographics, wealth, and caste interact are unobserved and could significantly influence whether the household decides to use family or hired women workers for rice cultivation. Even though we include these variables in our estimation, there might be unobserved factors related to these dimensions that we do not capture. Another important dimension that is omitted in our analysis is plot location, which could also guide whether family women work on the farm. It is plausible that the further away the plots are from the household's residence, the less likely family women are from working on them. Simultaneously, plot location could also influence agricultural productivity. In the absence of such data or a randomized assignment of the type of women's labor used on each plot, we test the sensitivity of our results to omission of variables using the methods proposed by Altonji, Elder, and Taber (2005) and Oster (2019). Their method is based on understanding the relative importance of omitted and the included variables and assessing the sensitivity of the estimated regression coefficients to omitted variables. Specifically, to examine the sensitivity of our results, we employ the strategy proposed by Oster (2019) and estimate the breakdown parameter, $\delta$, which measures the importance of selection on omitted factors relative to the observed, included control variables for our estimated results to not hold true. Mathematically, consider the following regression equation that we aim to estimate, where $X$ represents the household classification for the type of women's labor used, $Z_1$ is the vector of observed variables and $Z_2$ includes all the unobserved factors.

$$Y = \beta_0 + \beta_1 X + \alpha_1 Z_1 + \alpha_2 Z_2 + \epsilon \tag{12}$$

Based on Equation 12, the breakdown parameter, $\delta$, is measured as follows.

$$\delta = \frac{\frac{cov(X,\alpha_2 Z_2)}{var(\alpha_2 Z_2)}}{\frac{cov(X,\alpha_1 Z_1)}{var(\alpha_1 Z_1)}}$$

Here, the numerator captures selection on unobservables and the denominator measures selection on observables. To estimate $\delta$, we need the values of R-squared and $\beta$ under two conditions. As shown in Table 6, we need the value, $\beta_{short}$, which measures the regression coefficient we would obtain if we only regressed $X$ on $Y$ and did not include any control variables. For our study, $\beta_{short}$ is 0.15 and is obtained using a village-level fixed effect regression of whether the household uses only hired women's labor compared to family women's labor on output productivity. Next, the $\beta_{medium}$ value is the regression coefficient obtained using the regression of $Y$ on $X$ and $Z_1$, which equals 0.336, the coefficient estimated using village-level fixed effects and shown in Table 2, column (1). Next, we need the analogous $R^2_{short}$ and $R^2_{medium}$ estimated from the two models, which are 0.006 and 0.077. Using these parameters, Oster (2019) then allows us to estimate $\delta$ for different values of $R^2_{long}$, which is the $R^2$ value one would obtain if all unobservables were included in the regression model. As a rule of thumb, Oster (2019) suggests using a value of $1.3 * R^2_{medium}$ for $R^2_{long}$, which is approximately 0.1. Panel A in Table 6 estimates the $\delta$ parameter for a range of values of $R^2_{long}$, including 0.1, and tests the hypothesis that $\beta$ is equal to 0 (Panel A) and $\beta$ is greater than 0 (Panel B). As the values suggest, if we assume $R^2_{long}$ to be 0.1, then the estimated $\delta$ is 92 percent for the hypothesis that $\beta\,0$, which suggests that in our model, the selection on omitted factors would need to be almost as high as selection on observable, included factors for us reject the null hypothesis.

Next, we estimate the bias corrected $\beta$ coefficients that would result for using different values of the breakdown parameter, $\delta$, and holding $R^2_{long}$ constant (that is, $R^2_{long} = 0.1$). Panel C in Table 6 shows the sensitivity of our estimated coefficients for $\delta$ values ranging from $(-1, 1)$. A negative value of $\delta$ implies that the direction of the effect of the omitted factors is opposite to that of the observables. All the $\beta_{long}$ values obtained under different assumed values of $\delta$ are positive in magnitude. The lowest $\beta_{long}$ value we obtain is 0.24, which is very similar to the estimated regression coefficient obtained using the PDS LASSO estimation method, and the highest estimated value of $\beta_{long}$ is 0.47.

# 8 Conclusion

The paper highlights the existence of an agricultural productivity gap between households using family and hired women's labor, such that the forgone production value exceeds the costs associated with hiring women workers. We also find suggestive evidence that this gap is driven by a skill disparity between hired and family women workers. Our results are critical both because the magnitude of the productivity gap is substantial and because women's labor is a key input in rice cultivation, constituting approximately 37 percent of the total labor used in production. Yet, the policy implications based on these results need to be nuanced. Suggesting that investments target family women workers to improve their agricultural know-how would not only be simplistic but would also overlook the broader structural and social context in which these results arise (Doss, 2018). It is especially critical to situate these findings within the sectoral, household, and individual-level realities. First, the ongoing structural transformation in the rural sector, characterized by a shift of both men and women away from the agricultural sector, poses challenges. Whereas men are finding commensurate job opportunities in other sectors, women are leaving the rural labor force altogether (Raveendran and Kannan, 2012; Rangarajan, Kaul, and Seema, 2011). Unpaid family women workers constitute the biggest share of the women who have left the labor force. It is unlikely that this trend will reverse its course. As such, targeting investments to improve the skills of family women workers in agriculture may not offer the highest returns, especially if these women are likely to leave the agricultural sector and if there is limited inter-generational transmission of agricultural knowledge among women. In fact, it could be argued that the exit of family women workers, associated with lower agricultural productivity, might be agricultural beneficial for overall improvements in agricultural productivity.

The second reality relates to household welfare and influence of social and cultural norms on women's labor supply decisions. If the origin of family women's work on their farm and outside is rooted in social- and caste-based norms, then our results suggest that these agricultural productivity gaps are governed by more complex channels. They highlight that not only do these cultural and caste-based frictions impact women's work and productivity, but they also stand to influence the overall welfare of the household in terms of forgone agricultural productivity and income. Viewed from this perspective, it would be beneficial to target investments on mitigating these unobserved barriers that restrict women's labor supply decisions, in agriculture and beyond.

The third aspect involves considering the welfare and future work opportunities for both family and hired women workers. A problematic implication of our results arises if households perceive that women are better suited for unpaid household work rather than unpaid agricultural labor, and if agricultural innovation primarily aims to displace women's work. This shift also poses a risk to the welfare of hired women workers, particularly if a massive exit of women from the agricultural sector diminishes their bargaining power, leading to lower agricultural wage rates for women. Therefore, a more effective approach to investments would involve providing training and learning opportunities for family women to engage in off-farm jobs, alongside efforts to create rural off-farm work. More women-friendly job opportunities in the rural sector also strengthens the bargaining power of women working in the agricultural sector. Ultimately, any policy must be mindful of the intricate interplay of women's work in both agricultural and non-market domains, its interaction with unobserved norms, and its manifestation in household's agricultural production and consumption.

# References


Akresh, R. 2005. "Understanding Pareto Inefficient Intrahousehold Allocations." IZA Discussion Paper No. 1858.

Alesina, A., P. Giuliano, and N. Nunn. 2013. "On the Origins of Gender Roles: Women and the Plough." *The Quarterly Journal of Economics* 128:469–530.

Altonji, J.G., T.E. Elder, and C.R. Taber. 2005. "Selection on Observed and Unobserved Variables: Assessing the Effectiveness of Catholic Schools." *Journal of Political Economy* 113:151–184.

Belloni, A., V. Chernozhukov, and C. Hansen. 2014. "Inference on Treatment Effects After Selection Among High-Dimensional Controls." *Review of Economic Studies* 81:608–650.

Bindlish, V., R.E. Evenson, M. Gbetibouo, and B. Mundial. 1993. "Evaluation of T & V-Based Extension in Burkina Faso." Working paper, World Bank Washington, DC.

Blinder, A.S. 1973. "Wage Discrimination: Reduced Form and Structural Estimates." *Journal of Human Resources*, pp. 436–455.

Boserup, E. 1970. *Woman's Role in Economic Development*. New York: St. Martin's Press.

Cameron, C. 2022. "Machine Learning for Microeconometrics in 2022." Presentation slides.

Chen, M.A. 1989. "Women's Work in Indian Agriculture by Agro-Ecologic Zones: Meeting Needs of Landless and Land-Poor Women." *Economic and Political Weekly* 24:WS79–WS89.

Chernozhukov, V., D. Chetverikov, M. Demirer, E. Duflo, C.B. Hansen, W.K. Newey, and J.M. Robins. 2018. "Double/Debiased Machine Learning for Treatment and Structural Parameters." *Econometrics Journal* 21:C1–C68.

DiNardo, J., N.M. Fortin, and T. Lemieux. 1996. "Labor Market Institutions and the Distribution of Wages, 1973-1992: A Semiparametric Approach." *Econometrica* 64:1001–1044.

Doss, C.R. 2018. "Women and Agricultural Productivity: Reframing the Issues." *Development Policy Review* 36:35–50.

Eswaran, M., and A. Kotwal. 1985. "A Theory of Contractual Structure in Agriculture." *The American Economic Review* 75:352–367.

Eswaran, M., B. Ramaswami, and W. Wadhwa. 2013. "Status, Caste, and the Time Allocation of Women in Rural India." *Economic Development and Cultural Change* 61:311–333.

FAO, IFAD, and ILO. 2010. *Gender Dimensions of Agricultural and Rural Employment: Differentiated Pathways out of Poverty: Status, Trends and Gaps*. Food and Agriculture Organization of the United Nations, International Fund for Agricultural Development, and International Labour Office.

Farrell, M.H. 2015. "Robust inference on average treatment effects with possibly more covariates than observations." *Journal of Econometrics* 189:1–23.

Foster, A.D., and M.R. Rosenzweig. 1996. "Technical Change and Human-Capital Returns and Investments: Evidence From the Green Revolution." *The American Economic Review*, pp. 931–953.

Goldstein, M., and C. Udry. 2008. "The Profits of Power: Land Rights and Agricultural Investment in Ghana." *Journal of Political Economy* 116:981–1022, Publisher: University of Chicago.

IRRI. 1983. *Women in Rice Farming*. Manila: The International Rice Research Institute.

Jayachandran, S. 2015. "The Roots of Gender Inequality in Developing Countries." *Annual Review of Economics* 7:63–88.

Kazianga, H., and Z. Wahhaj. 2013. "Gender, Social Norms, and Household Production in Burkina Faso." *Economic Development and Cultural Change* 61:539–576.

Kitagawa, E.M. 1955. "Components of a Difference Between Two Rates." *Journal of the American Statistical Association* 50:1168–1194.

Luke, N., and K. Munshi. 2011. "Women as Agents of Change: Female Income and Mobility in India." *Journal of Development Economics* 94:1–17.

Mahajan, K. 2019. "Back to the Plough: Women Managers and Farm Productivity in India." *World Development* 124:104633.

Oaxaca, R. 1973. "Male-Female Wage Differentials in Urban Labor Markets." *International Economic Review*, pp. 693–709.

Oster, E. 2019. "Unobservable Selection and Coefficient Stability: Theory and Evidence." *Journal of Business & Economic Statistics* 37:187–204.

O'Sullivan, M., A. Rao, R. Banerjee, K. Gulati, and M. Vinez. 2014. "Levelling the Field: Improving Opportunities for Women Farmers in Africa." Working paper, The World Bank.

Paris, T., A. Singh, M. Hossain, and J. Luis. 2000. "Using Gender Analysis in Characterizing and Understanding Farm-Household Systems in Rainfed Lowland Rice Environments." *Characterizing and Understanding Rainfed Environments*, pp. 488.

Peterman, A., A. Quisumbing, J. Behrman, and E. Nkonya. 2011. "Understanding the Complexities Surrounding Gender Differences in Agricultural Productivity in Nigeria and Uganda." *The Journal of Development Studies* 47:1482–1509.

Quisumbing, A.R. 1996. "Male-Female Differences in Agricultural Productivity: Methodological Issues and Empirical Evidence." *World Development* 24:1579–1595.

Quisumbing, A.R., and L. Pandolfelli. 2010. “Promising Approaches to Address the Needs of Poor Female Farmers: Resources, Constraints, and Interventions.” *World Development* 38:581–592.

Rangarajan, C., P.I. Kaul, and Seema. 2011. “Where Is the Missing Labour Force?” *Economic and Political Weekly* 46:68–72.

Raveendran, G., and K.P. Kannan. 2012. “Counting and Profiling the Missing Labour Force.” *Economic & Political Weekly* 47:77–80.

Rosenzweig, M.R., and T.P. Schultz. 1982. “Market Opportunities, Genetic Endowments, and Intrafamily Resource Distribution: Child Survival in Rural India.” *The American Economic Review* 72:803–815.

Sahn, D.E., and H. Alderman. 1988. “The Effects of Human Capital on Wages, and the Determinants of Labor Supply in a Developing Country.” *Journal of Development Economics* 29:157–183.

Saradamoni, K. 1987. “Labour, Land and Rice Production: Women’s Involvement in Three States.” *Economic and Political Weekly*, pp. WS2–WS6.

Seymour, G. 2017. “Women’s Empowerment in Agriculture: Implications for Technical Efficiency in Rural Bangladesh.” *Agricultural Economics* 48:513–522.

Udry, C. 1996. “Gender, Agricultural Production, and the Theory of the Household.” *Journal of Political Economy* 104:1010–1046.

United States Department of Agriculture. 2023. “Country Summary - India.”

Wollburg, P., M. Tiberti, and A. Zezza. 2021. “Recall Length and Measurement Error in Agricultural Surveys.” *Food Policy* 100:102003.

Table 1: Summary Statistics, By Household's Use of Women's Labor

| | Only Family | | Both Family & Hired | | Only Hired | |
|---|---|---|---|---|---|---|
| | Mean | Std. Dev. | Difference | Std. Err. | Difference | Std. Err. |
| | (1) | | (2) | | (3) | |
| **Household Characteristics** | | | | | | |
| Household head's age (years) | 48.251 | (14.019) | -0.740 | (1.536) | 0.044 | (1.486) |
| Household head is male (%) | 97.861 | (14.507) | 1.266 | (0.993) | 2.139* | (1.140) |
| Household head's education (years) | 4.508 | (4.708) | 1.597*** | (0.456) | 4.673*** | (0.390) |
| Household is upper-caste (%)[a] | 8.556 | (28.047) | 1.488 | (4.367) | 46.056*** | (7.414) |
| Household is Hindu (%) | 95.722 | (20.291) | -0.525 | (2.584) | -2.364 | (3.327) |
| Household size | 5.984 | (2.693) | 0.047 | (0.314) | 0.023 | (0.332) |
| Wealth index[b] | -0.309 | (0.556) | 0.165** | (0.068) | 0.653*** | (0.093) |
| Eligible for loan (%) | 94.118 | (23.593) | -0.668 | (2.728) | 2.930 | (2.470) |
| Plots per household | 2.005 | (1.110) | 0.204* | (0.106) | 0.238* | (0.119) |
| Uses tractor (%)[c] | 70.053 | (45.925) | 14.663** | (6.303) | 21.459*** | (5.403) |
| Uses pump (%)[d] | 80.749 | (39.533) | 11.391* | (5.652) | 15.561*** | (5.523) |
| Uses pesticide (%) | 22.460 | (41.844) | 8.981 | (5.489) | 3.001 | (4.763) |
| Uses herbicide (%) | 4.278 | (20.291) | 5.329 | (4.126) | 1.626 | (2.496) |
| Uses fungicide (%) | 0.535 | (7.313) | 1.212 | (0.920) | -0.166 | (0.673) |
| Number of households | 187 | | 229 | | 271 | |
| **Plot Characteristics** | | | | | | |
| Productivity (kg/acre) | 1085.376 | (668.244) | 231.902** | (83.311) | 247.812** | (98.378) |
| Plot size (acres) | 0.548 | (0.855) | 0.086 | (0.105) | 0.439** | (0.167) |
| Years cultivated | 14.256 | (12.416) | 1.278 | (1.207) | 5.905*** | (1.264) |
| Plot is upland (%) | 12.267 | (32.849) | 2.358 | (5.799) | 1.878 | (4.660) |
| Plot is owned (%) | 49.333 | (50.062) | 8.967 | (6.243) | 35.864*** | (6.396) |
| Plot is share-cropped (%) | 36.533 | (48.217) | -9.458 | (6.121) | -27.652*** | (6.141) |
| Plot is rented-in (%) | 14.133 | (34.883) | 0.491 | (3.707) | -8.212** | (3.277) |
| Female labor days per acre | 31.020 | (24.941) | 6.469 | (3.928) | -12.27*** | (3.309) |
| Male labor days per acre | 57.013 | (43.485) | -10.322 | (6.112) | -9.654 | (5.689) |
| Irrigation hours per acre | 41.487 | (46.261) | 15.867** | (6.862) | 9.300* | (5.317) |
| Urea used (kgs/acre) | 151.639 | (248.867) | 20.324 | (38.032) | 44.010 | (36.560) |
| Uses non-urea fertiliser (%)[e] | 95.467 | (20.831) | 2.162 | (1.500) | -1.881 | (2.181) |
| Number of plots | 375 | | 506 | | 608 | |
| **Prices** | Mode | Mean | Std. Dev. | | | |
| Price of rice (INR/kg) | 10.000 | 45.952 | 155.074 | | | |
| Male transplanting daily wage (INR)[f] | 200.000 | 175.095 | 62.703 | | | |
| Female transplanting daily wage (INR)[g] | 125.000 | 133.678 | 49.765 | | | |
| Male weeding daily wage (INR)[h] | 100.000 | 117.920 | 52.334 | | | |
| Female weeding daily wage (INR)[i] | 150.000 | 115.024 | 45.212 | | | |

Notes: Standard deviation in parentheses for mean outcomes and clustered standard errors in parentheses for differences in mean outcomes. Statistical significance of differences is based on OLS regression with standard errors clustered at village-level. * p<0.1, ** p<0.05, and *** p<0.01.

[a] Households that do not identify as belonging to scheduled caste, schedule tribe, or other backward classes are classified as upper caste.

[b] Following variables were used in construction of the wealth index using factor analysis: ownership of cellphones, motorcycle, television units, cable television; expenditure on transport, education, and festival donations; ownership of diesel pump, rotavator, knapsack, and tractor; and the size of land owned (in acres).

[c] Use implies that the household either owns or rents tractors.

[d] Use implies that the household either owns or rents diesel or electric pumps.

[e] Non-urea fertilisers include potash, DAP (di-ammonium phosphate), SSP (single superphosphate), sulfur, and zinc.

[f] The summary statistics are calculated for 148 households reporting non-zero daily wages for male hired laborers in transplanting.

[g] The summary statistics are calculated for 376 households reporting non-zero daily wages for female hired laborers in transplanting.

[h] The summary statistics are calculated for 188 households reporting non-zero daily wages for female hired laborers in weeding.

[i] The summary statistics are calculated for 209 households reporting non-zero daily wages for female hired laborers in transplanting. The distribution of female weeding wages is bimodal, the two modes are 100 and 150.

Table 2: Agricultural Productivity and Type of Women's Labor Used

| Outcome variable: | Fixed Effects | | | PDS LASSO | | | AIPW LASSO | |
|---|---|---|---|---|---|---|---|---|
| Log of output productivity | (1) | (2) | (3) | (4)[a] | (5)[b] | (6)[c] | (7)[d] | (8)[e] |
| **Panel A: Estimates** | | | | | | | | |
| Only hired women | 0.336** | | 0.318*** | 0.261** | | 0.250*** | 0.377** | |
| | (0.124) | | (0.109) | (0.125) | | (0.095) | (0.185) | |
| Both family & hired women | | 0.179 | 0.169 | | 0.212** | 0.152* | | 0.439*** |
| | | (0.112) | (0.099) | | (0.096) | (0.085) | | (0.153) |
| **Panel B: Percent changes** | | | | | | | | |
| Only hired women | 40.000** | | 37.438*** | 29.813** | | 28.378*** | 45.841** | |
| Both family & hired women | | 19.573 | 18.364 | | 23.658** | 16.423* | | 55.144*** |
| Village fixed effects | Yes | Yes | Yes | Yes | Yes | Yes | No | No |
| Observations | 983 | 881 | 1489 | 983 | 881 | 1489 | 980 | 881 |
| R-squared | 0.077 | 0.062 | 0.070 | | | | | |
| Joint F-statistic (Both, Only hired = 0) | | | 4.950** | | | 14.382** | | |
| Difference (Only hired - Both) | | | 0.149** | | | 0.098* | | |
| Standard error | | | (0.059) | | | (0.057) | | |

Notes: Output productivity is measured as rice output (in kilograms) per acre. Robust standard errors shown in parentheses are clustered at village level. * p<0.1, ** p<0.05, and *** p<0.01. Columns (1), (2), and (3) show regression coefficients using a village-level fixed effect estimator; (4), (5), and (6) show regression coefficients using post-double selection LASSO estimator; and (7) and (8) show regression coefficients using the augmented inverse probability weighted LASSO estimator. Controls included in the fixed effects regressions: household head's sex, age, education; whether household is upper caste, non-Hindu; household's size, wealth index, eligibility for loan; whether household uses tractor, pump, pesticide, herbicide, and fungicide; plot size, years cultivated, tenancy, slope, and soil texture; female and male labor used per acre; hours of irrigation per acre; and urea and non-urea fertilizer used.

[a] From 650 control variables, 16 were selected by PDS LASSO.

[b] From 650 controls variables, 0 were selected by PDS LASSO.

[c] From 650 controls variables, 18 were selected by PDS LASSO.

[d] Village-level dummy variables were also included for selection using the AIPW LASSO procedure. From 677 controls variables, 13 were selected by AIPW LASSO. 3 observations which violated the overlap assumption were dropped.

[e] Village-level dummy variables were also included for selection using the AIPW LASSO procedure. From 677 controls variables, 6 were selected by AIPW LASSO.

Table 3: Placebo: Agricultural Productivity and Type of Men's Labor Used

| Outcome variable: Log of output productivity | Fixed Effects (1) | PDS LASSO (2)[a] | AIPW LASSO (3)[b] |
|---|---|---|---|
| Panel A: Estimates | | | |
| Use hired men's labor | 0.072 | 0.069 | 0.109 |
| | (0.078) | (0.069) | (0.105) |
| Panel B: Percent changes | | | |
| Use hired men's labor | 7.492 | 7.197 | 11.505 |
| Village fixed effects | Yes | Yes | No |
| Observations | 1489 | 1489 | 1489 |
| R-squared | 0.059 | | |

Notes: Output productivity is measured as rice output (in kilograms) per acre. Robust standard errors shown in parentheses are clustered at village level. * $p<0.1$, ** $p<0.05$, and *** $p<0.01$. Controls included in the fixed effects regressions: household head's sex, age, education; whether household is upper caste, non-Hindu; household's size, wealth index, eligibility for loan; whether household uses tractor, pump, pesticide, herbicide, and fungicide; plot size, years cultivated, tenancy, slope, and soil texture; female and male labor used per acre; hours of irrigation per acre; and urea and non-urea fertilizer used.

[a] Out of 650 control variables, 7 were selected by PDS LASSO.

[b] Village-level dummy variables were also included for selection using the AIPW LASSO procedure. Out of 677 control variables, 13 were selected by AIPW LASSO.

Table 4: Agricultural Productivity and Women's Labor Used, By Female Co-head's Wage Work

| Outcome variable: | Fixed Effects | | PDS LASSO | |
|---|---|---|---|---|
| Log of output productivity | Works Outside | Does not Work Outside | Works Outside | Does not Work Outside |
| | (1) | (2) | (3) | (4) |
| Panel A: Estimates | | | | |
| Only hired women | -0.303 | 0.431*** | 0.273 | 0.386** |
| | (0.248) | (0.149) | (0.389) | (0.164) |
| Panel B: Percent changes | | | | |
| Only hired women | -26.106 | 53.824** | 31.422 | 47.122** |
| Village fixed effects | Yes | Yes | Yes | Yes |
| Observations | 280 | 1042 | 280 | 1042 |
| R-squared | 0.467 | 0.089 | 0.247 | 0.112 |

Notes: Output productivity is measured as rice output (in kilograms) per acre. Working outside is a dummy variable capturing whether the female co-head reported engaging in paid work on other farms or off-farm. Robust standard errors shown in parentheses are clustered at village level. * p<0.1, ** for p<0.05, and *** for p<0.01. Controls included in the fixed effects regressions but not shown: household head's sex, age, education; whether household is upper caste, non-Hindu; household's size, wealth index, eligibility for loan; whether household uses tractor, pump, pesticide, herbicide, and fungicide; plot size, years cultivated, tenancy, slope, and soil texture; female and male labor used per acre; hours of irrigation per acre; and urea and non-urea fertilizer used. The set of covariates used in the PDS LASSO regressions are the same as those selected by the PDS LASSO procedure when implementing the regression for the whole sample, and as shown in column (4) of Table 2.

Table 5: Agricultural Productivity Gaps, By Share of Households Using Only Family Women's Labor in Village

| Outcome variable: Log of output productivity | Percentage of Households Using Only Family Women in Village ≤20% | ≤40% | ≤60% | ≤80% | ≤100% |
|---|---|---|---|---|---|
| | (1) | (2) | (3) | (4) | (5) |
| Panel A: Estimates | | | | | |
| Only hired women | 0.318 | 0.233** | 0.224** | 0.250** | 0.250*** |
| | (0.0.226) | (0.111) | (0.096) | (0.097) | (0.097) |
| Both hired and family women | 0.138 | 0.096 | 0.135 | 0.152* | 0.152* |
| | (0.158) | (0.113) | (0.089) | (0.087) | (0.087) |
| Panel B: Percent changes | | | | | |
| Only hired women | 37.478 | 26.191** | 25.146** | 28.372** | 28.378** |
| Both hired and family women | 14.830 | 10.065 | 14.478 | 16.420* | 16.423* |
| Village fixed effects | Yes | Yes | Yes | Yes | Yes |
| Observations | 604 | 1247 | 1455 | 1487 | 1489 |
| R-squared | 0.089 | 0.057 | 0.051 | 0.050 | 0.050 |
| Joint F-statistic (Both, Only hired = 0) | 1.953 | 3.525** | 2.927* | 3.423** | 3.428** |
| Difference (Only hired - Both) | 0.180* | 0.137** | 0.089 | 0.098 | 0.098 |
| Standard error (Only hired - Both) | (0.093) | (0.064) | (0.0.059) | (0.060) | (0.060) |

Notes: Output productivity is measured as rice output (in kilograms) per acre. Robust standard errors shown in parentheses are clustered at village level. * p<0.1, ** for p<0.05, and *** for p<0.01. The set of covariates used in the PDS LASSO regressions are the same as those selected by the PDS LASSO procedure when implementing the regression for the whole sample, and as shown in column (6) of Table 2.

Table 6: Sensitivity of Agricultural Productivity Gap to Omitted and Unobserved Factors

| | | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|
| Oster (2019) | | | | | | | | | | |
| $\beta_{short} = 0.151$;[a] $\beta_{medium} = 0.336$;[b] $R^2_{short} = 0.006$;[c] $R^2_{medium} = 0.077$[d] | | | | | | | | | | |
| Panel A: *Breakdown Value of $\delta$* | | | | | | | | | | |
| $H_0 : \beta_{Only\ hired\ women} \neq 0$ | | | | | | | | | | |
| $R^2_{long} =$ | 1 | 0.9 | 0.8 | 0.7 | 0.6 | 0.5 | 0.4 | 0.3 | 0.2 | 0.1 |
| $\delta(\%) =$ | 102.1 | 114.3 | 129.8 | 150.2 | 178.2 | 218.9 | 283.9 | 403.7 | 698.6 | 2590.1 |
| Panel B: *Breakdown Level of $\delta$:* | | | | | | | | | | |
| $H_0 : \beta_{Only\ hired\ women} > 0$ | | | | | | | | | | |
| $R^2_{long} =$ | 1 | 0.9 | 0.8 | 0.7 | 0.6 | 0.5 | 0.4 | 0.3 | 0.2 | 0.1 |
| $\delta(\%) =$ | 14.9 | 16.6 | 18.6 | 21.3 | 24.8 | 29.8 | 37.0 | 48.7 | 68.8 | 91.7 |
| Panel C: *Bounds of $\beta_{long}$ for $R^2_{long} = 0.1$*[e] | | | | | | | | | | |
| $\delta(\%) =$ | -0.99 | -0.8 | -0.6 | -0.4 | -0.2 | 0.2 | 0.4 | 0.6 | 0.8 | 0.99 |
| $\beta_{long} =$ | 0.24 | 0.26 | 0.28 | 0.30 | 0.32 | 0.36 | 0.38 | 0.41 | 0.44 | 0.47 |

Notes: The results presented are based on the the fixed effect regression of log of output productivity on the different labor-use household classification (controlling for plot-level, input use, and household-level factors) for the sample comprising households using only hired women and only family women. The regression results are presented in column (1) of Table 2.

[a] $\beta_{short}$ is the estimated coefficient associated with the "only hired women" household group using village-level fixed effect regression on log of productivity and without including any control variables. The sample comprises households using only family women and only hired women.

[b] $\beta_{medium}$ is the estimated coefficient associated with the "only hired women" household group using village-level fixed effect regression on log of productivity and includes plot-level, input use, and household-level characteristics. The sample comprises households using only family women and only hired women. The regression results are presented in column (1) of Table 2.

[c] $R^2_{short}$ is the R-squared for the fixed effect regression on log of productivity without any control variables, using the sample of only family women and only hired women households.

[d] $\beta_{medium}$ is the R-squared for the fixed effect regression on log of productivity with control variables, using the sample of only family women and only hired women households and as represented in column (1) of Table 2.

[e] Oster (2019) suggest $R^2_{long} = 1.3 * R^2_{medium}$ as a rule-of-thumb.

Figure 1: Types of Labor Used in Rice Cultivation

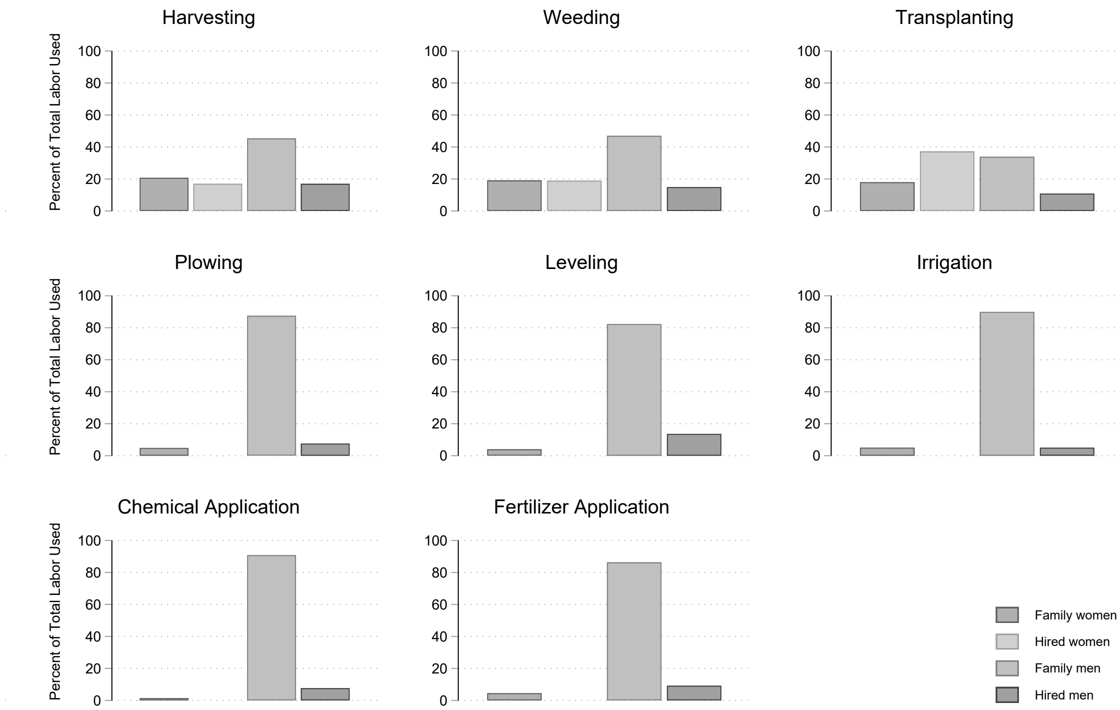


Notes: For each rice cultivation task, the share of each type of labor used is shown in the figures. The labor-use data were collected from the household head for each task. Very few households used any permanent labor (only 32 plots), so the shares do not reflect permanent labor used. N = 1489 plots.

Figure 2: Agricultural Productivity and Land Size, by Type of Women's Labor Used

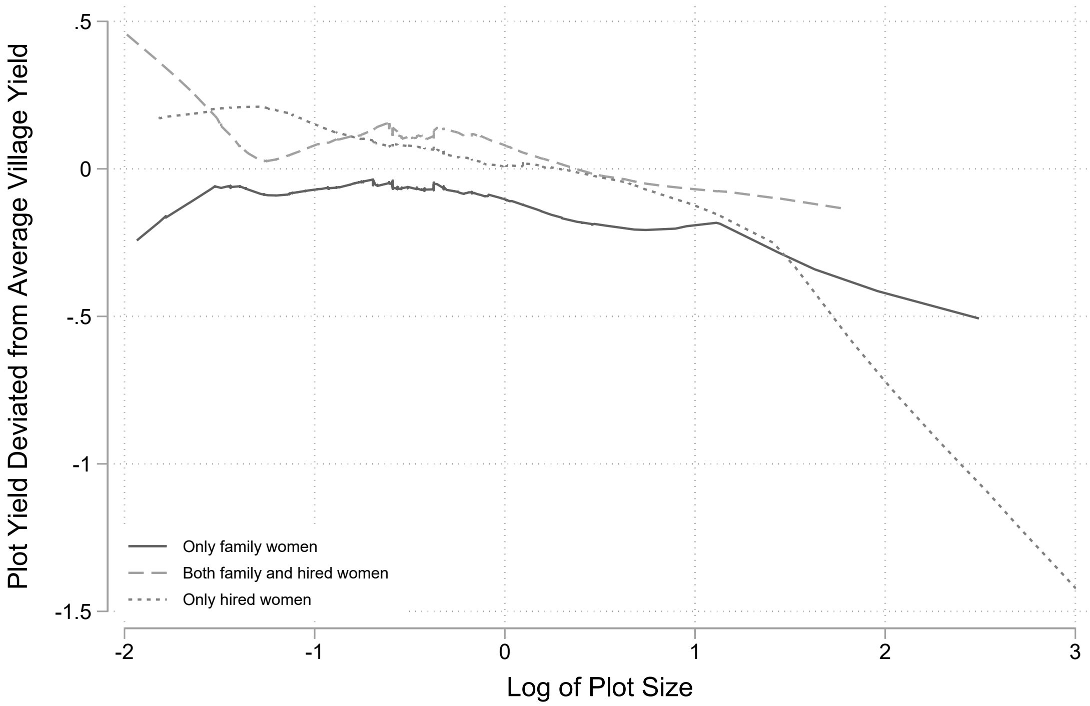


Notes: Plot yield is measured as the logarithm of rice output (in kilograms) per acre. Plot size is measured as logarithm of the plot size in acres. The figure uses locally estimated scatterplot smoothing (LOWESS) regressions to plot the relationship between deviation of plot yield and plot size, with Cleveland's (1979) tri-cube weighting function and running-line least squares smoothing. Bandwidth is set to 0.8. N = 1489 plots.

Figure 3: Actual and Counterfactual Agricultural Productivity Distributions

(a) Scenario 1: Counterfactual Distribution Based on Using Only Hired Women's Labor

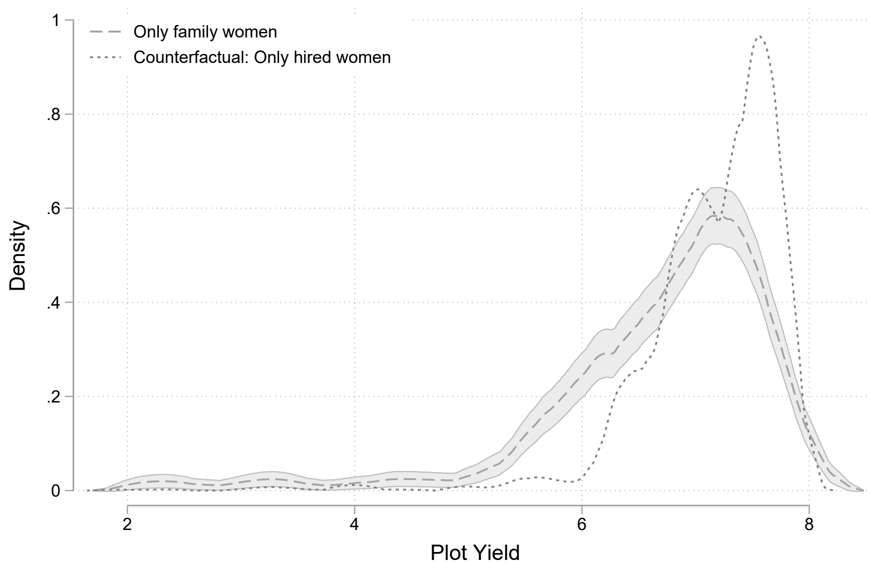


Notes: The figure compares the actual agricultural productivity distribution of households using only family women's labor with the counterfactual distribution that would prevail for households if they used only hired women's labor but had the same characteristics as households using family women's labor. N = 983 plots.

(b) Scenario 2: Counterfactual Distribution Based on Using Both Family and Hired Women's Labor

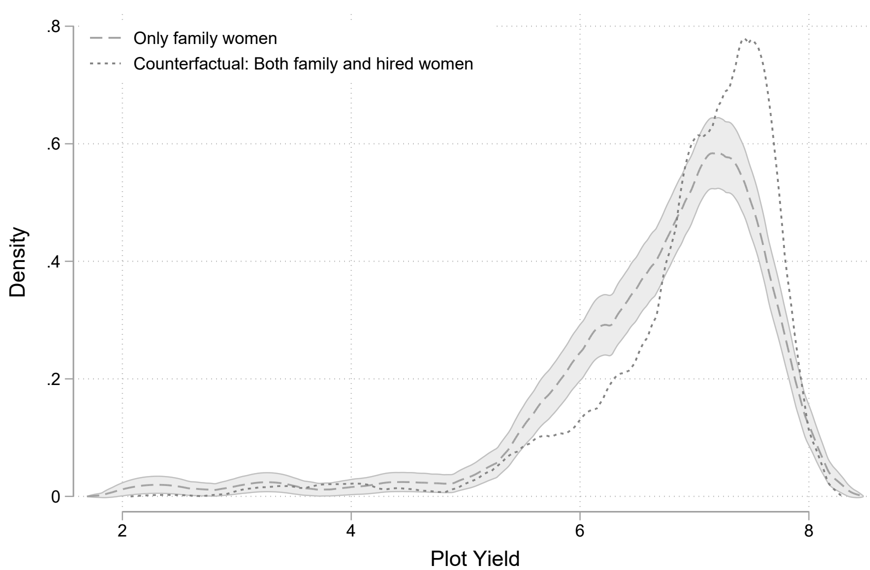


Notes: The figure compares the actual agricultural productivity distribution of households using only family women's labor with the counterfactual distribution that would prevail for households if they used both family and hired women's labor but had the same characteristics as households using family women's labor. N = 881 plots.

# Appendices

The standardized values of the following variables, their quadratic transformation, and interaction with one another were used in the PDS LASSO and AIPW LASSO selection procedure.

Table A1: Variables Used in PDS LASSO and AIPW LASSO Estimation

| **Household-level Variables** |
|---|
| Household head's age (years) |
| Household head is male (0/1) |
| Household head's education (years) |
| Household is upper-caste (0/1)[a] |
| Household is Hindu (0/1) |
| Household size |
| Wealth index[b] |
| Eligible for loan (0/1) |
| Uses tractor (0/1)[c] |
| Uses pump (0/1)[d] |
| Uses pesticide (0/1) |
| Uses herbicide (0/1) |
| Uses fungicide (0/1) |
| **Plot-level Variables** |
| Plot size (acres) |
| Years cultivated (years) |
| Plot is upland (0/1) |
| Tenancy type (Owned/Share-cropped/Rented-in) |
| Female labor days per acre |
| Male labor days per acre |
| Irrigation hours per acre |
| Soil type (Clayey/Sandy/Loamy/Clayey-Sandy/Sandy-Loamy/Black/Red) |
| Urea used per acre (kg/acre) |
| Uses non-urea fertiliser (0/1)[e] |

[a] Households that do not identify as belonging to scheduled caste, schedule tribe, or other backward classes are classified as upper caste.

[b] Following variables were used in construction of the wealth index using factor analysis: ownership of cellphones, motorcycle, television units, cable television; expenditure on transport, education, and festival donations; ownership of diesel pump, rotavator, knapsack, and tractor; and the size of land owned (in acres).

[c] Use implies that the household either owns or rents tractors.

[d] Use implies that the household either owns or rents diesel or electric pumps.

[e] Non-urea fertiliser include potash, DAP (di-ammonium phosphate), SSP (single super-phosphate), sulfur, and zinc.

Table A2: Agricultural Productivity and Women's Labor, Using Village-Level Fixed Effects

| Outcome variable: Log of output productivity | (1) | (2) | (3) |
|---|---|---|---|
| Only hired women | 0.336** | | 0.318*** |
| | (0.124) | | (0.109) |
| Both family and hired women | | 0.179 | 0.169 |
| | | (0.112) | (0.099) |
| Household head's age (years) | 0.001 | 0.000 | 0.001 |
| | (0.003) | (0.005) | (0.003) |
| Female-headed household (=1) | 0.310* | -0.177 | -0.023 |
| | (0.171) | (0.124) | (0.138) |
| Household head's education (years) | 0.007 | 0.013 | 0.011 |
| | (0.007) | (0.010) | (0.007) |
| Household is upper-caste (=1) | -0.099 | -0.083 | -0.078 |
| | (0.106) | (0.214) | (0.088) |
| Household is non-Hindu (=1) | -0.254** | 0.105 | -0.127 |
| | (0.113) | (0.316) | (0.151) |
| Household size | 0.001 | -0.005 | -0.007 |
| | (0.014) | (0.012) | (0.009) |
| Wealth index | -0.087** | -0.113* | -0.101*** |
| | (0.036) | (0.055) | (0.021) |
| Eligible for loan (=1) | -0.337*** | -0.146 | -0.235 |
| | (0.099) | (0.166) | (0.138) |
| Uses tractor (=1) | 0.002 | 0.036 | 0.036 |
| | (0.172) | (0.183) | (0.125) |
| Uses pump (=1) | -0.069 | -0.145 | -0.158 |
| | (0.194) | (0.183) | (0.136) |
| Uses pesticide (=1) | 0.200** | 0.121 | 0.124* |
| | (0.091) | (0.092) | (0.069) |
| Uses herbicide (=1) | -0.031 | 0.044 | 0.023 |
| | (0.126) | (0.165) | (0.111) |
| Uses fungicide (=1) | -0.652** | 0.212 | -0.151 |
| | (0.263) | (0.196) | (0.233) |
| Plot size (acres) | -0.082 | 0.014 | -0.029 |
| | (0.066) | (0.074) | (0.055) |
| Years cultivated | -0.046 | 0.025 | -0.022 |
| | (0.038) | (0.054) | (0.035) |
| Plot is upland (=1) | -0.122 | -0.049 | -0.085 |
| | (0.103) | (0.097) | (0.080) |
| Plot is share-cropped[a](=1) | 0.077 | -0.033 | 0.017 |
| | (0.133) | (0.117) | (0.083) |
| Plot is rented-in[a](=1) | -0.093 | -0.184 | -0.160* |
| | (0.143) | (0.119) | (0.089) |
| Soil: Clayey[b] | 0.275 | 0.883 | 0.630* |
| | (0.211) | (0.539) | (0.360) |
| Soil: Sandy[b] | 0.295 | 0.970 | 0.697* |
| | (0.222) | (0.572) | (0.387) |
| Soil: Loamy[b] | 0.270 | 0.814 | 0.601 |
| | (0.208) | (0.541) | (0.370) |
| Soil: Clayey-Sandy[b] | 0.312 | 0.709 | 0.535 |
| | (0.236) | (0.556) | (0.367) |
| Soil: Sandy-Loamy[b] | 0.333 | 1.039* | 0.717* |
| | (0.261) | (0.563) | (0.377) |
| Soil: Black[b] | 0.382* | 0.950* | 0.716* |
| | (0.211) | (0.541) | (0.353) |
| Female labor per acre | 0.060** | 0.013 | 0.060** |
| | (0.029) | (0.074) | (0.029) |
| Male labor per acre | 0.026 | 0.056 | 0.056 |
| | (0.046) | (0.052) | (0.037) |
| Irrigation hours per acre | 0.033 | 0.027 | 0.039 |
| | (0.048) | (0.039) | (0.034) |
| Urea used per acre | 0.049 | 0.054 | 0.052 |
| | (0.044) | (0.052) | (0.033) |
| Uses non-urea fertiliser[c](=1) | 0.044 | 0.313** | 0.095 |
| | (0.136) | (0.152) | (0.106) |
| Constant | 6.070*** | 5.137*** | 5.586*** |
| | (0.534) | (0.738) | (0.535) |
| Village-Level Fixed Effect | Yes | Yes | Yes |
| Observations | 983 | 881 | 1489 |
| R-squared | 0.077 | 0.062 | 0.070 |
| Joint F-statistic (Both, Only hired=0) | | | 4.950** |
| Difference (Only hired - Both) | | | 0.149** |
| Standard error | | | (0.059) |

Notes: Output productivity is measured as rice output (in kilograms) per acre. Robust standard errors shown in parentheses are clustered at village level. * p<0.1, ** p<0.05, and *** p<0.01.

[a] The omitted category is plots that are owned by the household.

[b] Non-urea fertilisers include: potash, DAP (di-ammonium phosphate), SSP (single superphosphate), sulfur and zinc.

[c] The omitted category is red soil.

Table A3: Placebo: Agricultural Productivity and Men's Labor, Using Village-Level Fixed Effects

| Outcome variable: Log of output productivity | |
|---|---|
| Uses hired men | 0.072 |
| | (0.078) |
| Household head's age (years) | 0.000 |
| | (0.003) |
| Female-headed household (=1) | -0.099 |
| | (0.092) |
| Household head's education (years) | 0.014** |
| | (0.006) |
| Household is upper-caste (=1) | -0.017 |
| | (0.090) |
| Household is non-Hindu (=1) | -0.061 |
| | (0.164) |
| Household size | -0.009 |
| | (0.009) |
| Wealth index[c] | -0.093*** |
| | (0.022) |
| Eligible for loan (=1) | -0.244* |
| | (0.131) |
| Uses tractor (=1) | 0.067 |
| | (0.125) |
| Uses pump (=1) | -0.156 |
| | (0.141) |
| Uses pesticide (=1) | 0.124* |
| | (0.070) |
| Uses herbicide (=1) | 0.011 |
| | (0.102) |
| Uses fungicide (=1) | -0.230 |
| | (0.217) |
| Plot size (acres) | -0.030 |
| | (0.060) |
| Years cultivated | -0.020 |
| | (0.036) |
| Plot is upland (=1) | -0.086 |
| | (0.079) |
| Plot is share-cropped[a] (=1) | 0.068 |
| | (0.087) |
| Plot is rented-in[a] (=1) | -0.143 |
| | (0.085) |
| Soil: Clayey[b] | 0.629* |
| | (0.345) |
| Soil: Sandy[b] | 0.696* |
| | (0.376) |
| Soil: Loamy[b] | 0.605 |
| | (0.359) |
| Soil: Clayey-Sandy[b] | 0.548 |
| | (0.355) |
| Soil: Sandy-Loamy[b] | 0.740* |
| | (0.365) |
| Soil: Black[b] | 0.734** |
| | (0.343) |
| Female labor per acre | 0.043 |
| | (0.028) |
| Male labor per acre | 0.054 |
| | (0.037) |
| Irrigation hours per acre | 0.038 |
| | (0.032) |
| Urea (kgs) used per acre | 0.054 |
| | (0.032) |
| Uses non-urea fertilisers[c] (=1) | 0.078 |
| | (0.107) |
| Constant | 5.712*** |
| | (0.536) |
| Village-Level Fixed Effects | Yes |
| Observations | 1489 |
| R-squared | 0.059 |

Notes: Output productivity is measured as rice output (in kilograms) per acre. Robust standard errors shown in parentheses are clustered at village level. * p<0.1, ** p<0.05, and *** p<0.01.

[a] The omitted category is plots that are owned by the household.

[b] Non-urea fertilisers include: potash, DAP (di-ammonium phosphate), SSP (single superphosphate), sulfur and zinc.

[c] The omitted category is red soil.

Table A4: Robustness: Agricultural Productivity and Type of Women's Labor Used for Sub-Sample of Male Headed Households

| Outcome variable: | Fixed Effects | | | PDS LASSO | | | AIPW LASSO | |
|---|---|---|---|---|---|---|---|---|
| Log of output productivity | (1) | (2) | (3) | (4)[a] | (5)[b] | (6)[c] | (7)[d] | (8)[e] |
| Panel A: Estimates | | | | | | | | |
| Only hired women | 0.337** | | 0.326*** | 0.255** | | 0.266*** | 0.379** | |
| | (0.125) | | (0.110) | (0.127) | | (0.096) | (0.186) | |
| Both family & hired women | | 0.184 | 0.176* | | 0.216** | 0.159* | | 0.443*** |
| | | (0.112) | (0.100) | | (0.097) | (0.085) | | (0.155) |
| Panel B: Percent changes | | | | | | | | |
| Only hired women | 40.094** | | 38.514*** | 29.009** | | 30.538*** | 46.142** | |
| Both family & hired women | | 20.246 | 19.271* | | 24.161** | 17.202* | | 55.794*** |
| Village fixed effects | Yes | Yes | Yes | Yes | Yes | Yes | No | No |
| Observations | 978 | 870 | 1478 | 978 | 870 | 1478 | 978 | 870 |
| R-squared | 0.078 | 0.063 | 0.071 | | | | | |
| Average productivity | 1239.932 | 1220.723 | 1267.029 | 1239.932 | 1220.723 | 1267.029 | 1239.932 | 1220.723 |
| Joint F-statistic (Both, Only hired = 0) | | | 5.059** | | | 15.828** | | |
| Difference (Only hired - Both) | | | 0.150** | | | 0.108* | | |
| Standard error | | | (0.064) | | | (0.062) | | |

Notes: Output productivity is measured as rice output (in kilograms) per acre. Robust standard errors shown in parentheses are clustered at village level. * p<0.1, ** for p<0.05, and *** for p<0.01. Columns (1), (2), and (3) show regression coefficients using a village-level fixed effect estimator; (4), (5), and (6) show regression coefficients using post-double selection LASSO estimator; and (7) and (8) show regression coefficients using the augmented inverse probability weighted LASSO estimator. Controls included in the fixed effects regressions: household head's sex, age, education; whether household is upper caste, non-Hindu; household's size, wealth index, eligibility for loan; whether household uses tractor, pump, pesticide, herbicide, and fungicide; plot size, years cultivated, tenancy, slope, and soil texture; female and male labor used per acre; hours of irrigation per acre; and urea and non-urea fertilizer used.

[a] From 650 control variables, 15 were selected by PDS LASSO.

[b] From 650 controls variables, 1 were selected by PDS LASSO.

[c] From 650 controls variables, 15 were selected by PDS LASSO.

[d] Village-level dummy variables were also included for selection using the AIPW LASSO procedure. Out of 677 controls variables, 13 were selected by AIPW LASSO.

[e] Village-level dummy variables were also included for selection using the AIPW LASSO procedure. From 677 controls variables, 6 were selected by AIPW LASSO.

Table A5: Robustness: Agricultural Productivity and Type of Women's Labor Used for Sub-Sample of Lower-Caste Households

| Outcome variable: | Fixed Effects | | | PDS LASSO | | | AIPW LASSO | |
|---|---|---|---|---|---|---|---|---|
| Log of output productivity | (1) | (2) | (3) | (4)[a] | (5)[b] | (6)[c] | (7)[d] | (8)[e] |
| Panel A: Estimates | | | | | | | | |
| Only hired women | 0.277* | | 0.287** | 0.295* | | 0.284** | 0.403** | |
| | (0.143) | | (0.127) | (0.159) | | (0.127) | (0.184) | |
| Both family & hired women | | 0.172* | 0.159 | | 0.210** | 0.181* | | 0.458*** |
| | | (0.098) | (0.096) | | (0.097) | (0.097) | | (0.166) |
| Panel B: Percent changes | | | | | | | | |
| Only hired women | 31.949* | | 33.187** | 34.304* | | 32.880** | 49.647** | |
| Both family & hired women | | 18.793* | 17.190 | | 23.418** | 19.808* | | 58.017*** |
| Village fixed effects | Yes | Yes | Yes | Yes | Yes | Yes | No | No |
| Observations | 600 | 797 | 1054 | 600 | 797 | 1054 | 600 | 797 |
| R-squared | 0.075 | 0.076 | 0.073 | | | | | |
| Average productivity | 1198.436 | 1208.272 | 1244.636 | 1198.436 | 1208.272 | 1244.636 | 1198.436 | 1208.272 |
| Joint F-statistic (Both, Only hired = 0) | | | 2.575* | | | 9.996* | | |
| Difference (Only hired - Both) | | | 0.128 | | | 0.104 | | |
| Standard error | | | (0.080) | | | (0.075) | | |

Notes: Output productivity is measured as rice output (in kilograms) per acre. Robust standard errors shown in parentheses are clustered at village level. * p<0.1, ** for p<0.05, and *** for p<0.01. Columns (1), (2), and (3) show regression coefficients using a village-level fixed effect estimator; (4), (5), and (6) show regression coefficients using post-double selection LASSO estimator; and (7) and (8) show regression coefficients using the augmented inverse probability weighted LASSO estimator. Controls included in the fixed effects regressions: household head's sex, age, education; whether household is upper caste, non-Hindu; household's size, wealth index, eligibility for loan; whether household uses tractor, pump, pesticide, herbicide, and fungicide; plot size, years cultivated, tenancy, slope, and soil texture; female and male labor used per acre; hours of irrigation per acre; and urea and non-urea fertilizer used.

[a] From 650 control variables, 3 were selected by PDS LASSO.

[b] From 650 controls variables, 1 were selected by PDS LASSO.

[c] From 650 controls variables, 8 were selected by PDS LASSO.

[d] Village-level dummy variables were also included for selection using the AIPW LASSO procedure. From 677 controls variables, 5 were selected by AIPW LASSO.

[e] Village-level dummy variables were also included for selection using the AIPW LASSO procedure. From 677 controls variables, 3 were selected by AIPW LASSO.

Table A6: Agricultural Productivity and Women's Labor, Using Household Fixed Effects

| Outcome variable: Log of output productivity | (1) |
|---|---|
| No family women | 0.191** |
| | (0.091) |
| Plot size (acres) | -0.034 |
| | (0.087) |
| Years cultivated | -0.129 |
| | (0.086) |
| Plot is upland (=1) | 0.300*** |
| | (0.103) |
| Plot is share-cropped[a](=1) | 0.240 |
| | (0.166) |
| Plot is rented-in[a](=1) | -0.091 |
| | (0.170) |
| Soil: Clayey[b] | 0.020 |
| | (0.483) |
| Soil: Sandy[b] | 0.265 |
| | (0.500) |
| Soil: Loamy[b] | 0.192 |
| | (0.469) |
| Soil: Clayey-Sandy[b] | -0.047 |
| | (0.464) |
| Soil: Sandy-Loamy[b] | -0.093 |
| | (0.453) |
| Soil: Black[b] | 0.227 |
| | (0.448) |
| Female labor days per acre | -0.043 |
| | (0.108) |
| Male labor days per acre | 0.150* |
| | (0.085) |
| Irrigation hours per acre | 0.026 |
| | (0.028) |
| Urea (kg) per acre | 0.102** |
| | (0.039) |
| Uses non-urea fertiliser[c](=1) | 0.119 |
| | (0.154) |
| Constant | 5.706*** |
| | (0.712) |
| Fixed Effect | Household-level |
| Observations | 1489 |
| R-squared | 0.057 |

Notes: Output productivity is measured as rice output (in kilograms) per acre. Robust standard errors shown in parentheses are clustered at village level. * p<0.1, ** p<0.05, and *** p<0.01. The "No family women" variable indicates that no family women works on the plot in the household. Within households using only family women's labor, from 375 plot observations, only 12 plots did not have any family women working on them. Among households using both family and hired women's labor, from 506 plot observations, only 44 plots did not have family women working on them. All plots in households using only hired women's labor do not have any family women working on them. Thus the reported estimates should be interpreted with caution due to very little within-household variation in the type of labor used.

[a] The omitted category is plots that are owned by the household.

[b] Non-urea fertilisers include: potash, DAP (di-ammonium phosphate), SSP (single superphosphate), sulfur and zinc.

[c] The omitted category is red soil.

Figure A7: Placebo: Actual and Counterfactual Agricultural Productivity Distributions Based on Using Men's Labor Also

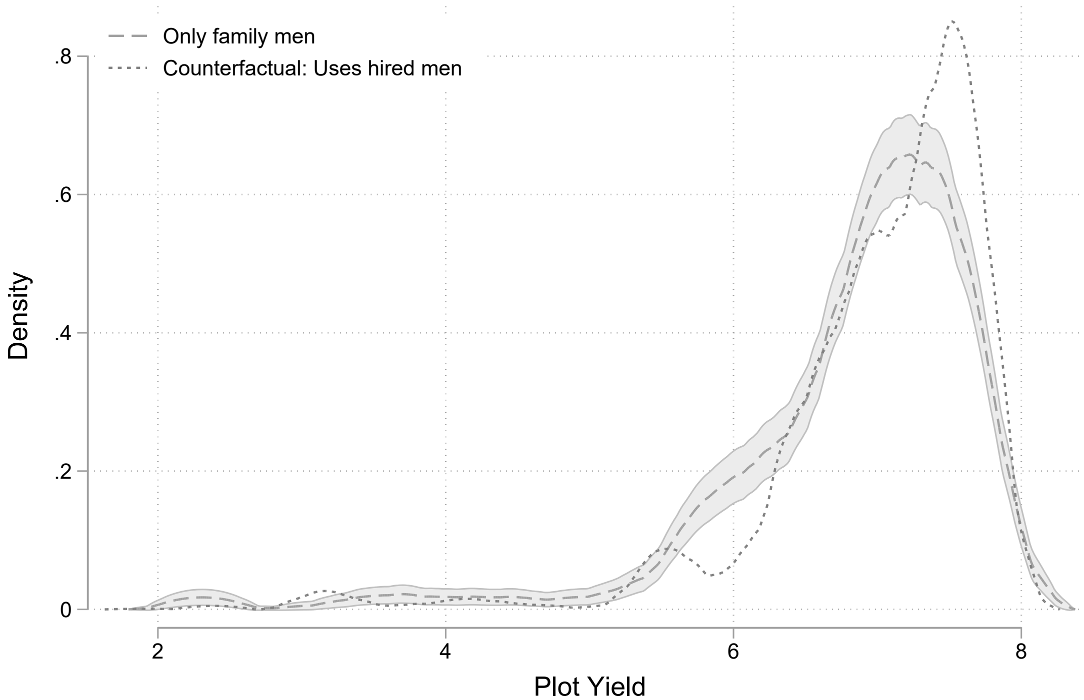


Notes: The figure compares the actual agricultural productivity distribution of households using only family men's labor with the counterfactual distribution that would prevail for households if they also used hired men's labor but had the same characteristics as households using only family men's labor. N = 1489 plots.