Literacy and Female Status in Green Revolution India

Juan Pablo Rud*

March 26, 2009

Abstract

I investigate the effect of the Green Revolution in India on rural literacy and women's status using a difference in differences strategy by cohort and intensity of adoption in a panel of 254 districts. I find robust evidence of a substantial increase in literacy associated with the adoption of the new agricultural technology. The effect is present for the treated cohorts only – i.e. young and rural– but not for older rural and young urban cohorts. I also provide results that help reconcile the conflicting anthropological evidence on women's status in Green Revolution India: even though, on average, labour, educational and survival outcomes worsen over time for rural Indian women, a greater adoption of the new technology is shown to mitigate this trend.

1 Introduction

The introduction of High Yield Varieties (HYV) during the Green Revolution in India in the mid-1960s significantly increased food production in a short period of time. In just five years, all India foodgrain production increased by almost 50% and average yields increased by around 30%. However, the heterogeneous effects associated with the introduction of the new technology generate conflicting evidence in terms of whether it brought sizeable benefits to the rural population. In particular, many authors (see, for example, Sharma and Dak (1989) for a review) argue that only relatively well-off farmers and large landowners benefited, while the many poor and small farmers, tenants or agricultural labourers might have experienced marginal benefits at best. As the new technology was intensive in irrigation, fertilizers and mechanization, a successful introduction of the new seeds required not just an initial

*Deparment of Economics and DESTIN, London School of Economics. I thank Fernando Aragon, Robin Burgess, Esteban Jaimovich, Guy Michaels, Steve Redding and Jon Temple for their useful comments and suggestions.

level of prosperity, but also an adequate level of education. That means that even if states' governments provided subsidised seeds, fertilizers and machinery to farmers–relaxing the financial constraint to adopt the new technology–educated farmers would still be better prepared to process information and enjoy the advantages of technical change. In that context, Foster and Rosenzweig (1996) point out that more educated farmers became aware earlier or simply were more able to manage the new technology. They show that the likelihood of adoption of HYV seeds was significantly larger for households with at least an adult that had completed primary education, controlling for measures of wealth such as land size. Additionally, after showing that profits were greater for educated households, controlling for other inputs such as access to irrigation and machinery, the authors found that enrollment rates in primary schools increased in areas that experienced a larger increase in agricultural yields. In short, because returns to primary education increased, private investments in schooling also increased. But returns to education increased at a much faster pace for areas or individuals that started off relatively better, sending them into a path of more growth and higher incomes. The relationship between technical change and inequality raises the concern as to whether the increased returns to education at the micro level generated effects of any significance at the aggregate level, considering that a majority of the rural population might have been excluded from the process of technical change, where inputs were complementary to education[1].

Alternatively, the unequal distribution of benefits associated to the Green Revolution might also come as a consequence of dimensions of inequality other than wealth or education. A particularly salient dimension in the literature is that of gender (Sridhar, 2004). Most notably, issues on female status in rural India had been underlined by the problem of "missing women" –i.e. a statistically important difference in the number of boys and girls–by Sen (1990). Additionally, a low participation in markets (labour, credit, etc) could also exclude them from the adoption of the new technology and improvements in well-being that could result from it (see, for example FAO (1997)). Anthropological evidence of the effect of the introduction of HYV seeds on female status is somewhat conflicting. Some authors, such as Rosenzweig and Schultz (1982) and UNDP (2003), for example, show that women employment actually increased in areas with characteristics associated with HYV adoption, sometimes at a faster pace than men's. Others, in contrast–such as Boserup (1990), Harriss (1989), Mencher (1988) and Mazumdar and Sharma (1990)–have argued that the technological change had a pro-male bias, because of the complementarity between mechanization and male labour. If that is the case, then the advent of the new seeds

[1] For example, Arora et al.(1989) show that between 1971 and 1981, the average operational holding increased from 2.89 hectares to 3.79 hectares. In particular, small and marginal holdings (up to one hectare) were 37.63% in 1971 and 19.21% in 1981, while large holdings (between 3 and 10 hectares) increased from 26.11% to 37.01% of all holdings. The authors suggest that "small and marginal holders were leasing out or selling their lands, (...), since farming on a small scale had become an uneconomic proposition."

should have been associated with a relative reduction of women in the workforce or, at least, in the agricultural sector, deepening the rift between genders. More generally, displacement from the agricultural sector would also imply that relative income and returns to education for women were lower, and that educational improvements should have been lower for girls than for boys in rural areas. Qian (2008), for example, shows that in rural China relative improvements in rural female income were associated with improvements in educational attainment for all children, while an improvement in males' relative income benefited boys' education only. In the context of Green Revolution India, Foster and Rosenzweig (1996) found a result that suggests a technological pro-male bias: investments in education were significantly larger for boys than for girls, on average. However, this result does not account for relative changes in female working status in different HYV intensive areas (e.g. by type of crop), that could show heterogeneous gender effects of technology adoption on educational attainment.

In this paper, I test the relationship between HYV adoption and aggregate literacy and female status at the district level in India, using census data for 254 Indian districts in 13 major Indian states for the census years of 1961 (pre-Green Revolution), 1971 (less than five years after its start) and 1981. The use of pre-Green Revolution data, census information and actual data on HYV adoption (as opposed to agricultural yields) are the main differences with Foster and Rosenzweig's paper. The data for this paper are described in Appendix A and presented in Table 1. There is a great variation in the intensity of HYV adoption across districts and across time. The mean proportion of cropped area under HYV seeds was around 10% in 1971, only four years after the introduction of the new seeds and rose to above 25% ten years later. Still, some districts never introduced the new technology, while others did so very quickly and deeply. Rural literacy overall also increased over time, for both the male and female populations, even though the levels of literacy were much higher for men than for women, and the gap in literacy rates evaluated at the mean increased from 24 to 27 percentage points by 1981, compared to the pre-Green Revolution gap. This suggests a deteriorating gender gap in rural areas that is also showing in the participation of rural women in the labour force. It is striking to see that, on average, almost 80% of the women in rural areas in 1971 were not in the labour force, jumping from 66% in 1961. At first glance, the last two described patterns in the data would seem to support the idea that there was an increase in the gender gap in rural India after the start of the Green Revolution. However, I show below that districts intensive in HYV were not generating this process.

The main specification I use is similar to the one developed by Duflo (2001) to test the impact of school construction in Indonesia in schooling outcomes, i.e. a difference in difference regression where I compare the change in literacy for rural cohorts that were in primary schooling age during the Green Revolution in HYV intensive districts, with respect to the same cohort in 1961 and in districts less intensive in HYV. The identifying assumption is that there are no omitted time-varying district characteristics correlated with HYV adoption and rural literacy that would generate spurious results. To deal with this concern in the identification strategy, in addition to controlling for district and state varying characteristics, I use two control groups: older rural cohorts that would capture pre-Green Revolution trends in literacy and the same cohorts in urban areas to capture contemporaneous district-wide trends in literacy. In short, I look both at the average treatment effect on the treated (where the treatment is districts' HYV intensity) and at the average treatment effect on the untreated (older rural and urban cohorts) to test whether the treatment is picking up other unobserved phenomena that could have improved literacy even in the absence of the Green Revolution.

I find robust evidence that the increase in adoption of HYV seeds is associated with substantial increases in literacy levels. Literacy increased in all districts on average, but a district at the mean of the HYV adoption distribution would have produced an extra increase in rural literacy of around 2 percentage points per cohort. The reduced form estimation means that I cannot tell the different channels driving the increase in rural literacy apart. It could be the case that local governments built schools or hired more teachers in HYV intensive areas, increasing the supply of education or that higher returns to education have increased demand. I find evidence that when controlling for state supply of education, HYV adoption predicts greater literacy suggesting that, at least, the demand channel is at work. My results are consistent with the idea that a greater complementarity between economic activity and educational levels is an effective way of creating demand-side incentives for education. This can be of particular importance in rural areas, where the persistence of traditional farming might be associated with low returns to education.

Furthermore, I find no evidence of a Green Revolution-related increase in the gender gap but quite the opposite: a greater intensity in HYV is associated with a decrease in the literacy gap between boys and girls and with more women in the labour market, in the agricultural sector in particular. The interesting story the results show is that in rural areas there is an average increase in the literacy gender gap and a decrease in female labour participation. However, a more intensive level of HYV adoption mitigates these negative effects, without offsetting them completely. I found a similar result for the survival of girls relative to boys aged between 0 and 4 in rural areas. These findings could shed some light on the conflicting evidence regarding gender issues and the Green Revolution: anthropological research looking at a small number of districts over time might conclude that the rural gender gap has increased with and because of the Green Revolution, without noting that the trend in most Indian districts was the same and even worse in areas with lower or no HYV adoption.

In addition, I provide evidence that the effects of the Green Revolution on literacy levels are present irrespectively of the gender specificity of crops. Despite the fact that female participation in the labour market and in agricultural activities were more prevalent in areas that were relatively intensive in rice (rather than in wheat) before the introduction of HYV seeds, literacy for both boys and girls improved with the adoption of both high yield varieties of wheat and of rice. However, I do find that female literacy improved relatively more than male literacy in areas that adopted more HYV of rice. That suggest that the relative impact of HYV adoption of a given crop on literacy was stronger for the gender whose labour is relatively more intensive in that crop. This result is consistent with the evidence found by Qian (2008) for rural China, where reforms that improved relatively more a crop intensive in female labour had positive effects on female status.

Results are robust to the use of geographic characteristics to predict HYV adoption in an instrumental variables (IV) setting. Some concerns related to the endogeneity of HYV adoption remain and are not dealt with by the cohort analysis and placebo tests. For example, it could be the case that improvements in younger rural cohorts' literacy are driving the modernisation of the agricultural sector or that districts that are undergoing an unobserved cultural transformation are those who simultaneously improve female status and technology adoption in rural areas. To adress this issue, I use the time-varying effects of time invariant geographic characteristics to predict divergence in HYV adoption, as in Rud (2009). In particular, the geological thickness of the watertable (needed for irrigation) and the low relief of the terrain are salient characteristics of technology-adopting districts, and their variation is enough to predict improvements in literacy and female status in the second stage of the IV regression.

The paper is organised as follows. The next section analyses the effect of the Green Revolution on literacy. Section 3 looks at different indicators of female status, according to participation in labour markets, literacy, survival rate of girls and by looking at the gender-specificity of crops. In section 4 I show IV results and section 5 concludes.

2 Literacy and the Green Revolution

2.1 Difference in differences analysis by cohort

As a first step in identifying whether the introduction of HYV brought about an increase in literacy, in this section I carry out a difference in difference analysis of the evolution of literacy before and after the start of the Green Revolution. The census data provide information on rural literacy for different age groups at different points in time which can be used to distinguish those cohorts that went through the educational age during the Green Revolution years, from those cohorts of individuals who did not. Additionally, I use information on the proportion of area that was cropped with HYV seeds at the district level to distinguish districts that were intensive in HYV (above the median) and those districts that were not.

The baseline year I use is 1961, five to six years before the first HYV seeds were introduced. For 1971, I use as the treated cohort those children between the age of

5 and 14–i.e. that were at most 10 years of age when the Green Revolution started. Table 2, Panel A shows that areas which were above the HYV adoption median in 1971 had significantly higher levels of rural literacy in 1961 (3.5 percentage points). In line with the idea that education was an important input for the successful adoption of the new seeds, this first difference suggests that the Green Revolution was introduced more rapidly in areas with better initial education. By comparing the difference in 1961 between high and low HYV intensity areas to the difference in 1971, column (3) shows that the wedge not just persists, but experiences a significant increase, both statistically (at 1% level) and in magnitude (the difference is more than 65% larger). It is interesting to note that literacy increases over time for both groups of districts. This approach controls for group characteristics with time-invariant effects (by comparing the same group at two points in time) and common time shocks. However, this approach relies on the identifying assumption that there are no time-varying district specific effects that are correlated with literacy other than by HYV adoption, i.e. that rural literacy would not have improved more in these districts if the HYV seeds were never introduced.

To test whether the pattern of increase in literacy is driven by systematic differences across both groups of districts, I carry out two placebo experiments within the simple difference in difference framework. In the first one, I look at the literacy levels in the cohort of rural individuals aged 15 to 24–who were not exposed to primary education under the Green Revolution. Given the breakdown in age groups provided by the census data, this measure is very conservative, since many children in this group were 11 or 12 when the Green Revolution started and might have been able to improve their educational efforts at that time. In the second control experiment, I run a similar difference in difference approach by looking at the cohort of urban children aged between 5 and 14. In both cases, if regional trends were at work and were confounded with the adoption of HYV, I should observe a similar pattern of literacy increases in either older rural cohorts, in urban children or in both. The second part of Panel A and Panel B in Table 2 show that there is no systematic difference between high and low HYV areas in literacy changes for rural individuals aged 15 to 24 and for urban children aged between 5 and 14. For the former, there is a significant difference between high and low HYV intensity regions in 1961 (supporting the idea mentioned previously, that HYV was more successful in better educated areas), but this difference does not change after the start of the Green Revolution (see Column (6)). This suggests that the results obtained for the younger rural cohort were not picking up pre-existing trends in rural literacy in areas that ended up adopting more HYV seeds[2]. A difference in difference across cohorts in a given year shows that in 1971, the difference between high and low areas in the treated cohort is significantly higher than the difference for the older cohort. Interestingly, the cross-cohort difference in difference was negative in 1961 and turned positive after the start of the

[2]Unreported results show that the same pattern of no difference in high and low HYV areas holds for the group of people older than 25.

Green Revolution. Overall, Column (7) shows that the change between 1961 and 1971 is 2.1 percentage points (significant at the 1% level).

For the cohort of urban children aged 5 to 14, Panel B shows three interesting features. First, there is no difference in literacy before the Green Revolution in urban areas for regions high and low in HYV adoption. Second, urban literacy has increased between 1961 and 1971 by similar magnitudes to the ones I found in rural areas (around 5 or 6 percentage points). This suggests that increases in literacy were not an exclusively rural phenomenon at the time. Finally, however, this increase over time is very similar in both high and low HYV districts. In short, for children aged 5 to 14, I don't find any evidence that the increase in literacy in urban areas was more prevalent in districts intensive in HYV. In terms of the identifying assumptions, it is encouraging to find that changes that show in treated cohorts are present in non-treated cohorts.

In Table 3, I carry out a similar exercise, this time comparing data for 1961 and 1981. This time, the cohort that grew up after the Green Revolution started includes those aged between 5 and 24 in 1981, and I compare them with the same group in 1961. Districts are again sorted according to whether the proportion of HYV adoption is above or below the median, measured in 1981. Panel A shows that the difference in rural literacy for those aged between 5 and 24 between both regions is 2.9 percentage points greater (at the 1% significance level) in 1981 than in 1961.

As in the previous table, I look for systematic differences across regions by testing whether the same difference in difference approach can explain the evolution of two groups that are not treated under my identifying assumptions: i.e. rural populations aged over 25 and urban populations aged below 25. Results in Panels A and B in Table 3 show that the increasing levels of rural literacy found for the treated group are not present for the other two control groups. The difference in literacy between high and low HYV areas for the group that had their educational years after the start of the Green Revolution is significantly greater than the difference in literacy for the same group in 1961, the older age group in 1981 and the same age group in urban areas[3].

This section has provided evidence that rural literacy significantly improved in HYV intensive areas. The claim that the difference in difference results are identifying causal effects seems more robust when results for untreated cohorts in rural and urban areas do not follow the same pattern, alleviating concerns that the results are driven by time-varying regional characteristics. In the following section, I extend the analysis to use a continuous measure of HYV adoption and to control for district-specific and time-varying characteristics that could be correlated with HYV adoption and literacy and might be biasing the difference in difference results in Tables 2 and 3.

[3] The results hold when comparing the 15 to 24 age group in 1971 (not treated) and 1981 (treated), showing that the change in literacy in HYV intensive districts between 1971 and 1981 is on average 1.5 percentage points greater (significant at 1% level) than in low HYV areas. Older rural and urban cohorts don't show any significant difference.

2.2 Effect of HYV adoption by cohort

In this section, I turn to a regression framework to introduce a continuous measure of HYV adoption and control for district characteristics that might be driving the surge in rural literacy. I first consider the average effect on literacy of the Green Revolution by running an OLS regression of the following form

$$L^{cl}_{dst} = \alpha_{ds} + \beta_t + \gamma HYV_{dst} + \delta \mathbf{X}_{dst} + \varepsilon_{dst} \tag{1}$$

where L^{cl}_{dst} is the proportion of literate people in district d, state s and year t for a cohort c in location l (i.e. rural or urban), and HYV_{dst} is the time varying proportion of cropped area cultivated with HYV seeds at the district level. I also include district fixed effects to control for time-invariant district characteristics and year dummies to control for average changes in literacy between census years. I also include a set of district controls $\mathbf{X}_{dst}$ that capture demographic changes (e.g. proportion of female population, proportion of rural population, log of population) and economic characteristics (e.g. proportion of rural workers in the manufacturing sector, log of roads per square kilometre) that could be correlated with both HYV adoption and literacy. To account for autocorrelation of the error term, I cluster standard errors in all regressions at the district level.

Even though the previous section has shown that the successful adoption of HYV seeds was more prevalent in areas with higher literacy, these estimates are not capturing this reverse causality. The fact that the variable HYV_{dst} is only positive and different from 0 from 1971, implies that district fixed effects are absorbing the differences in initial levels of literacy across districts and that the estimated coefficient γ only captures the additional improvement in literacy associated with increases in HYV. Note that the estimation is a reduced form regression, without being able to delineate specific mechanisms, such as an increase in demand for education or a Green Revolution-induced increase in the supply of schools in high HYV areas. As suggested by Foster and Rosenzweig (1996), both likely happened to some extent. However, the main econometric concern would be that, in the absence of technological change, these districts would have experienced supply or demand shocks at that time specific to rural areas. In the absence of a clean counterfactual, the use of a continuous measure of HYV adoption allows me to compare districts with different treatment intensity over time, controlling for district fixed effects and time-varying characteristics. Additionally, running similar specifications for non-treated cohorts in rural and urban areas allows me to check whether the literacy changes were happening only at the right time and in the expected areas.

Panel A in Table 4 shows results for three cohorts, with and without district controls in each case. In columns (1) and (2), increases in HYV adoption are positively associated with greater literacy rates for the group of rural children aged 5 to 14. The inclusion of district controls reduces the magnitude of the coefficient, but not its significance. Additionally, the year dummies are positive and significant, capturing an average improvement in rural literacy among the youngest cohort. The magnitude is sizeable: an increase in one standard deviation in HYV results in around 1.4 percentage points increase in rural literacy for children aged 5 to 14. Taken at the average district in the sample, that means having almost 5400 more literate children in that cohort. For the whole sample, this is equivalent to an increase in the number of literate children aged 5 to 14 by an extra 1.4 million, as a result of the introduction of HYV seeds.

It is important to note that the cohort of children aged 5 to 14 were of educational age in both HYV years (1971 and 1981). That means that I expect the effect of HYV on this cohort to be stronger than for the cohorts analysed in the subsequent columns. In columns (3) and (4), I run regression (1) using the cohort of individuals in rural areas aged 15 to 24. Given that the Green Revolution started in 1966/7, I expect many people in the 1971 cohort to be beyond educational age, probably in the labour market already. However, all individuals in the 1981 cohort were born (and then were eligible to receive primary education) after the start of the Green Revolution. Because the cohort in 1971 includes many people that were not treated, I expect the effect of HYV on literacy to be smaller than for the youngest cohort. Results in Columns (3) and (4) show that this is the case, with coefficients remaining strongly significant. When pooling both cohorts together in columns (5) and (6) and running on rural literacy for individuals aged 5 to 24, I include all treated individuals (aged 5 to 14 in 1971 and aged 5 to 24 in 1981) and some that I expect not to be treated (aged 15 to 24 in 1971). The estimates are positive and strongly significant (at the 1% level) and, as expected, the magnitude of the coefficients is smaller than the estimates in columns (1) and (2), but greater than in columns (3) and (4).

Panel B in Table 4 looks at the correlation between literacy and some district controls to shed some light on some of the drivers of rural literacy. The proportion of population that is female is strongly correlated with higher levels of literacy while more rural districts don't seem to perform worse. However, more densely populated areas do have lower literacy rates across cohorts. As expected, measures of rural economic development and infrastructure–such as the proportion of rural workers in manufacturing jobs and road density–are both positively correlated with higher levels of literacy.

In Table 5, I reproduce regression (1) but this time using the control groups, as in the previous section. As in Table 4, I include regressions with and without district controls. In columns (1) and (2) the left hand side variable is the rate of literacy for the rural cohort of individual that was of educational age before the start of the Green Revolution, i.e. those aged 25 or more. As expected, the coefficient on HYV adoption is not significantly different from 0. The year dummies are positive and significant, suggesting an increasing improvement in rural literacy across generations that predates independence, but that was not associated with the Green Revolution. Columns (3) to (6) use the groups aged 5 to 14 and 5 to 24 in urban areas to see whether districts' intensity in HYV is correlated with urban literacy among the cohorts that were of educational age after the introduction of HYV seeds. A positive and significant coefficient would imply that HYV adoption was correlated with omitted district developments which were driving literacy up in both rural and urban areas and that the results are capturing a spurious correlation between HYV adoption and rural literacy. The failure to produce statistically significant coefficients for urban literacy alleviates these concerns somewhat.

Next, I combine the baseline OLS regression and the difference in difference approach to look at the additional effect per year of HYV adoption with respect to 1961 in a regression of the following form:

$$L_{dst}^{cl} = \alpha_{ds} + \beta_t + \gamma_{71} HYV_{ds71} + \gamma_{81} HYV_{ds81} + \delta \mathbf{X}_{dst} + \varepsilon_{dst} \quad (2)$$

where γ_{71} and γ_{81} respectively capture the differential increase in literacy associated with higher HYV intensity in 1971 and 1981 with respect to 1961. The previous OLS specification provided information on the average effect of HYV on literacy. Equation (2) allows me to deal with at least three concerns. The first, related to the magnitude of the effect of the Green Revolution on rural literacy, is to deal with the possibility that in order to get sizeable improvements in literacy, the adoption of HYV seeds should increase beyond the bounds of the data at a point in time. I can check that by having a measure of the effect per year. A second concern is that a cohort that was treated shows significant results because the same cohort in another year is pulling the average up. For example, if all literacy improvements for rural children aged 5 to 14 are significant only in 1981, and insignificant or negative in 1971, this could raise doubts as to whether the improvements came from other sources than HYV adoption, which only became active many years after the start of the Green Revolution. Similarly, this specification shows whether a cohort that was not supposed to be treated (e.g. individuals aged 15 to 24 in 1971) produces significant results. A third concern is that the effects observed in the OLS regression were short-lived, thus capturing a pre-Green Revolution trend in districts that ended up adopting more HYV, rather than its effects. If that is the case, the differential increase in rural literacy associated with HYV intensity with respect to 1961 could be greater in 1971 than in 1981 (where HYV adoption was greater). In short, since HYV adoption increased steadily over time, finding an unbalanced distribution of the effects of HYV on rural literacy over time might raise some concerns about the described channel: specifically, it might be revealing that the measure of HYV in a given year is capturing some other omitted time-varying district-specific effects that are driving rural literacy up.

Table 6 shows results for the same three cohorts analysed earlier, with and without district controls. Columns (1) and (2) show that for rural children aged 5 to 14, a greater proportion of HYV seeds is associated with an additional increase in rural literacy over and above the average increase captured by the year dummies. As expected, the effect holds both in 1971 and 1981, meaning that the effects of the Green Revolution on literacy are present from the beginning. In terms of magnitude, a district at the mean of HYV adoption in 1971 has a 1.1 percentage points extra increase in rural literacy (on top of the average increase of 3.7 percentage points) with respect to 1961. When comparing 1981 with 1961, a district at the mean of the adoption distribution has increased between 2.3 and 3.2 percentage points more than the average increase of 9 percentage points. Note that even though the point estimates are similar, the fact that HYV was more prevalent in 1981 increases the mean effect and alleviates the concern that the effects on literacy were not distributed over time in a similar way to HYV adoption. Columns (3) and (4) show interesting results regarding the cohort of rural individuals aged between 15 and 24. Since most of the individuals in this cohort were not supposed to be treated in 1971, I expect the effects to be insignificant or small at best. This is what the results show: I find no effect in the specification without district controls, and I find a coefficient significant at the 10% level only when I control for district characteristics. If present, the magnitude of the effect is low (0.5 percentage points, less than half the effect found for the younger cohort). This contrasts with the positive, sizeable (around 2 percentage points) and significant (at the 1% level) effect found for this cohort in 1981. These results are good news in terms of the identifying assumptions, since the effects on rural literacy seem to come not just from districts that were more intensive in HYV, but also from the cohorts that were supposed to be affected by the Green Revolution. Columns (5) and (6) show that the results hold when pooling together the two cohorts, i.e. individuals aged 5 to 24, and that the magnitudes are very similar to those found in the first two columns.

The same exercise is done for control groups, namely rural people that were in schooling age before the Green Revolution and younger urban cohorts. A concern with respect to the OLS results for these control groups might be (again) that OLS is averaging out some opposing effects that might happen in HYV-intensive districts over time, revealing some unobserved phenomena that have affected different cohorts at different points in time. If during the first years of the Green Revolution, adopting districts were drawing resources from urban areas to serve rural areas, then we should expect intensity in HYV to be associated with a drop–or less than average increases– in urban literacy. Alternatively, resources generated by HYV in the late stages might have been used to fund a district-wide schooling programme that would affect both urban and rural children. If stories of this sort happened over time, it could be the case that an OLS estimation would show an insignificant effect of HYV on urban areas, but a difference in differences approach might show a positive effect for 1971, and a negative effect for 1981, for example. Results in Table 7 alleviate the concern that a more complex non-linear story was generating insignificant results in urban or older rural cohorts. Even though the year dummies capture increasing literacy in all groups, there is no strong evidence that intensity in HYV was associated with heterogeneous effects for the untreated population at any point in time.

Finally, I check whether the estimates are confounding the effects of other phenomena that could be generating a spurious correlation between HYV adoption and rural literacy[4]. The first alternative story is one related to migration: the initial difference in difference tables showed that HYV adoption was more successful in areas with greater levels of literacy. It might well be the case that, as a result, areas suitable for HYV adoption were attracting more literate farmers. If that were the case, the increase in average rural literacy would not be the consequence of increased returns to education generating incentives to younger generations to get more education, but instead the consequence of a selection of better educated migrants arriving in the HYV-compatible districts. Column (1) in Table 8 controls for the interaction of the proportion of rural migrants in the district with year dummies. The estimates for HYV adoption in 1971 and 1981 remain positive, significant and similar in magnitude, suggesting that HYV was not capturing the effects of migration.

Another concern is that HYV adoption was correlated with the proportion of population belonging to Scheduled Castes (SC) or Scheduled Tribes (ST). Dhanagare (1989) points out that the benefits for these groups were smaller, in particular because it was poorer farmers that on average benefited less from the introduction of agricultural change. The story I am testing here is one where the evolution of literacy in a district is linked to the economic and social composition of the population, rather than to HYV. The presence of a poorer or more disenfranchised population could be correlated both with less HYV adoption and with a deteriorating trajectory for rural literacy. Columns (2) and (3) check whether the effect of HYV vanishes when controlling for the proportion of population belonging to either group, interacted with year dummies. Since the proportion of people belonging to either group in a district tends to be very stable, the interaction with year dummies will capture the time evolution of rural literacy for a level of SC/ST population. The effects of HYV remain positive and significant, even though the magnitude drops a little, in particular when the proportion of population belonging to Scheduled Tribes increases[5].

Finally, the increase in literacy might come as a consequence of improved provision of education supplied by the state government and unrelated to the adoption of the new technology. In column (4), I include as a control the log number of schools per capita at the state level[6]. In Column (5), I also include state time-trends to check whether changes in literacy were actually driven by state-wide policy trajectories. In both cases, results on the interaction terms are positive and significant, and the magnitude of the coefficients does not drop much. These specifications reduce concerns that literacy improvements were not driven by HYV adoption but by other district or state phenomena that would have improved rural literacy in the absence of the

[4] I present results for the cohort of children aged 5 to 14, even though the following results hold for the other cohorts as well.

[5] Interestingly, the unreported coefficients on the interaction of the ST measure with year dummies are negative and significant. This implies that the upward bias comes from a negative correlation between successful adoption of HYV and more population classified as Scheduled Tribe, as suggested by the anecdotal evidence.

[6] I alternate this measure with the log of expenditure in primary education per capita and the log of primary teachers per capita, without major changes in the results.

Green Revolution.

3 Gender, literacy and the Green Revolution

In this section, I analyse whether the improved literacy associated with HYV intensity had heterogeneous effects by gender. It is generally argued that women lost out in the process of technical change associated with the Green Revolution in India, either by being displaced from wage earning opportunities or by receiving even lower wages than before (see FAO (1997), for example). Theoretically, if the new technology was a substitute rather than a complement to female labour, I should observe fewer women employed in rural areas or, if employed, that they moved away from the agricultural sector. Unfortunately, the data available does not allow me to investigate whether women with particular characteristics (e.g. landowners or from wealthier households) benefited more. But the data do allow me to investigate if claims that female labour in general–and in agriculture in particular–was displaced in areas intensive in HYV, hold on aggregate. The existing evidence is not conclusive about whether women were displaced by the introduction of technical change (as suggested by Harriss (1989) or Mencher (1988)) or whether female employment actually increased (as in UNDP (2003), for example). Kaur and Sharma (1989) suggest that there are heterogeneous effects from mechanization for women based on socioeconomic characteristics. They find that the overall effect is a displacement of female labour in rural areas in two districts in the state of Haryana.

In terms of literacy outcomes, there are two effects that should be considered: one is linked to parents' demand for their children's education when a more profitable technology has greater returns to education than the traditional agricultural technique. The second one is linked to intra-household allocation of power and resources: there is evidence that when resources shift towards women, all children benefit and when resources shift towards men, only boys do, if anyone (see Qian (2008) and Rosenzweig and Schultz (1982) for China and India, respectively). It follows that increasing returns to education stemming from HYV adoption should be a force inducing greater literacy levels for all children; even though if resources have shifted towards males, the effects on boys should be greater than on girls. Moreover, in the presence of strong gender bias literacy changes for girls could go either way, i.e. I could find that the interactions of HYV and year dummies in equation (2) are negative or insignificant when the explained variable is rural female literacy for a treated cohort.

To test the direction of effects, I will first check labour market outcomes for women, in absolute terms and relative to men. I subsequently test whether improvements in literacy were gender specific by looking separately at the effects on gender cohorts and the gap between male and female literacy and by analysing the effects by gender-intensity of the two major crops during the Green Revolution (i.e. rice and wheat). I finally check whether HYV can be associated with changes on another measure of female status, namely, girl's relative survival rate.

3.1 Labour market and gender

The effect of the Green Revolution on female labour could theoretically go either way. A simple mechanism for increasing female labour is that the new technology increased female wages and raised the opportunity cost of not working. However, if the technology was intensive in inputs that were not complementary to female labour, then women's wages would have dropped. Unless rural women were in a backward-bending part of the labour supply schedule, a drop in wages should have reduced the supply of female labour. However, other mechanisms might have been at play if the decision to work was made at the household level or depended on other characteristics, such as land ownership. For example, if HYV adoption increased the income of male household members, women might have stopped working altogether– or moved away from the agricultural sector. On the other hand, the adoption of the new seeds by the household might only have been possible if women raised some extra cash to buy seeds, fertilizers, etc. The district data do not provide enough detail to tell these alternative mechanisms apart. However, a reduced form regression of employment outcomes on HYV adoption might be informative in terms of labour market outcomes, particularly since the data set includes districts with little or no HYV adoption. To that end, I use information on the proportion of women not working and, among those who did work, the proportion employed in agriculture as a proportion of total rural females and working females, and relative to men.

In column (1), Table 9, the year dummies in the OLS regression show that the proportion of women that were not part of the working force in rural areas was greater in 1971 than in 1961, but not in 1981. However, a greater level of HYV adoption was associated with a reduction in the proportion of women that did not participate in the labour market. Note that the proportion of HYV adoption that offsets completely the effect captured by the year dummy should be above 1. That means that HYV adoption could at best mitigate average changes in female employment. This suggests that following a district intensive in HYV before and after the introduction of the new technology would show that female employment dropped but would not reveal what is apparent from results in column (1): that the effect was actually lower than it would have been in the absence of the new technology.

Column (2) looks at the interaction of HYV adoption with year dummies and finds that a greater part of the effect of HYV adoption in column (1) happened by 1971. However, even though the 1981 dummy is still insignificant, intensity in HYV was associated with a slight increase in female participation in the labour force. The magnitude of the coefficients and the distribution of HYV adoption per year suggest that, with the exception of a small number of districts, female participation in labour markets dropped significantly between 1961 and 1971, but the effect was increasingly mitigated by a greater adoption of HYV seeds. By 1981, the proportion of women not working fell again, but the increase in female labour (when compared to 1961 levels) came from districts with greater adoption of HYV only.

In columns (3) to (6) I look at the participation of women in the agricultural sector.

When measured as a proportion of the total female population, columns (3) and (4) show that intensity in HYV attracted more women into the agricultural sector. The HYV coefficient in column (3) is positive and significant and column (4) shows that the effect was stronger in 1971. It is apparent again that HYV adoption mitigated the decrease in the proportion of women in the agricultural sector, as captured by the 1971 dummy. Not including a proper counterfactual (i.e. evolution of districts with less HYV or without HYV altogether) might lead to the conclusion that HYV was associated with a displacement of female force from agricultural labour. Columns (5) and (6) look at the proportion of rural women in agriculture, conditional on being in the labour force, to check the sectorial allocation of female workers. Over time, there seems to have been almost no significant movement towards or away from agriculture among female workers. However, an increase in HYV adoption was strongly associated with women working in the agricultural sector. This time, the effects are not just strongly significant but also within the bounds of the data: the increase in the proportion of female workers in the agricultural sector was 2.3 percentage points higher for a district at the mean of the HYV adoption distribution in 1971, and 5.6 percentage points higher in 1981 than in a district with no HYV.

Finally, columns (7) and (8) take a look at the ratio female labour - male labour in the agricultural sector. Results show that HYV intensity is associated with relatively more female labour used, even though the trend was decreasing. These effects only show just after the start of the Green Revolution.

3.2 Literacy

Table 9 provides evidence that the Green Revolution not just attracted women to the labour market but to the agricultural sector, in particular. To test whether the increase in female employment was associated with increasing returns to education for females, in the absence of information on wages for female agricultural workers, I can only run a reduced form regression and check whether literacy among females increased with HYV. If an increase in the labour market gap between males and females resulting from the Green Revolution was not captured by Table 9, but was present and affected the returns to education, then it should be reflected in the evolution of the literacy rates by gender.

To analyse whether the literacy results found in previous sections were gender specific, I run regression (1) again, this time for men and women separately. Information on literacy rates per age groups and per sex allows me to test the hypothesis that the Green Revolution induced a gender bias in educational outcomes. Columns (1) and (2), Table 10 present results for the OLS estimation of HYV adoption on literacy for rural individuals aged 5 to 24 for males and females, respectively. An increase in 1 percentage point in the proportion of land cropped with HYV seeds was associated with an increase in rural literacy of 0.06 percentage points for males and 0.10 percentage points for females. As the average effect obtained in the OLS estimate for this cohort was 0.08, it follows that there were heterogeneous effects in rural literacy associated with HYV adoption, but, if anything, in a pro-female direction. To check whether this differential impact in literacy came from gender specific trends, I will again use the two control groups (rural population older than 25 and the population aged 5 to 24 in urban areas). The average effect in these groups was not significantly different from zero in the OLS estimation. But it could be the case that female literacy was following a different trajectory in rural and urban areas at the district level (through a catch up effect, for example, since rural male literacy was more than two times greater than female literacy, on average, or through a district-wide programme that also reached urban girls). So it could be that column (2) was simply reflecting these effects. Columns (3) and (4) show that for the older rural cohort HYV was associated with changes in literacy that are insignificantly different from zero. Similarly, results for the urban cohort in columns (5) and (6) are also insignificant. In both cases, the results on HYV in the first two columns do not seem to be picking up a gender-specific rural or generational effect.

As pointed out before, looking at the average effect of HYV on rural literacy might be hiding some heterogeneous gender effects behind the positive and significant estimates within cohorts. The steady increase in HYV adoption over time should be reflected in the year specific effects in order to reduce concerns that the measure of HYV adoption is correlated with year and region-specific shocks. To test this, I ran a set of difference in difference regressions as in equation (2), but this time for male and female cohorts, to see whether the gender effect on literacy found in the OLS specification followed the deepening of the Green Revolution. For the cohort of rural individuals aged between 5 and 24, the point estimates confirm that the effect of HYV was stronger in 1981 than it was in 1971, for both men and women. In particular, intensity with HYV was associated with a larger and increasing effect on female literacy. In a district with a mean level of HYV adoption, the increase in rural literacy was similar for men and women in 1971 (i.e. 0.8 and 1 percentage points, respectively), but larger for females in 1981 (1.6 percentage points for males, 2.6 for females)[7]. When looking at the control groups, columns (3) to (6) show no results that could suggest that omitted characteristics at the district levels were driving the general improvement in literacy or the relatively improved performance of female literacy in rural or urban areas.

The positive and significant coefficients for year dummies across generations and locations show that literacy did improve for men and women. Moreover, improvements above the average in areas with higher HYV intensity only happened for the cohorts that were supposed to be affected by the change, i.e. individuals that were educated after the start of the Green Revolution in rural areas. Additionally, the improvements in literacy seem to be evenly distributed between men and women; if anything, with a tilt towards larger benefits for girls. In the next set of regressions,

[7] Unreported estimates show a very similar pattern for the cohort of rural children between 5 and 14, suggesting that the choice of cohort is not relevant in identifying gender effects.

I check whether these differences are statistically significant, by running a regression of the form:

$$\Delta L^{cl}_{dst} = \alpha_{ds} + \beta_t + \gamma_{71} HYV_{ds71} + \gamma_{81} HYV_{ds81} + \delta \mathbf{X}_{dst} + \varepsilon_{dst} \quad (3)$$

where ΔL^{cl}_{dst} is the literacy gap between men and women of cohort c in location l for a district d in state s and year t. The difference should cancel out all district characteristics that had a similar effect on male and female literacy (e.g. teacher availability or quality), but would not deal with characteristics that had a differential impact on literacy, even if these characteristics were fixed (e.g. cultural gender bias). To account for this possibility, I use district fixed effects that would take care of gender-varying effects of district fixed characteristics. Similarly, year dummies will capture the change in the gender gap with respect to 1961. Additionally, I replace the left hand side variable with the difference between the gender gap in rural areas with the gender gap in urban areas, to account for the possibility that the adoption of HYV seeds might have been correlated with some unobserved district-wide changes in gender attitudes.

The positive and significant estimates for the year dummies in column (1) in Table 12 suggest that the rural literacy gap increased between males and females aged between 5 and 24. However, the gender gap increase was mitigated in districts with more HYV adoption, i.e. the increase in the literacy gap was lower, the greater the incidence of HYV seeds. Only in a few cases was the level of HYV adoption large enough to offset the overall increase in the literacy gap. Column (2) shows the gender gap in literacy for the same cohort in urban areas. The year dummies show that, on average, the gap decreased with respect to 1961. If anything, in 1971 the gap decreased slightly less in districts more intensive in HYV, even though the magnitude of the coefficients is small and a reversion in the decrease captured by the year dummy is beyond the bounds of the HYV adoption distribution. Columns (3) and (4) show results for the rural and urban cohort of people aged over 25 and prove interesting: as in the previous two columns, the gender gap decreased in urban areas and increased in rural areas. Most notably, districts intensive in HYV showed a reduction in the gender literacy gap. This would be bad news in terms of the identification strategy, since there was an effect in HYV-intensive districts in a cohort that should not have been affected. However, column (4) shows that a similar phenomena was happening in urban areas, suggesting that districts intensive in HYV were closing the gender literacy gap in rural and urban areas at the district level, probably capturing pre-Green Revolution trajectories. The next four regressions use the difference between rural and urban gender gaps as the explained variable, with the objective of absorbing common trends in gender gap literacy in urban and rural areas. In columns (5) and (7) I run an OLS regression using HYV intensity as the explanatory variable and in columns (6) and (8), I run regression (3). Column (5) shows that the gender gap in rural areas increased over time vis-a-vis urban areas (as shown by the point estimates of the year dummies), but the coefficient on HYV adoption is negative and significant.

This means that the difference between the rural and urban gender gaps decreased with the Green Revolution. Column (7) shows that this was not the case for the older cohort which grew up before the Green Revolution started. Columns (6) and (8) show the same pattern: the relative increase in the gender gap in rural areas was mitigated in areas with greater HYV adoption for the cohort that was treated, i.e. of educational age after the start of the Green Revolution, and not for the older cohorts.

In conclusion, results in this section suggest not just that the intensity of the Green Revolution was associated with an increase in literacy in rural areas for both men and women, but that the effect was significantly stronger for women. The evidence suggests that HYV adoption actually mitigated the increasing gap between male and female rural literacy. This observation is at odds with the widespread belief that the Green Revolution worsened the status of women, at least in terms of access to education (see FAO report on "Women and Food Security" (1997) and Mazumdar and Sharma (1989)). Since the gender literacy gap in rural areas seems to have worsened over time, looking at the evolution of districts intensive in HYV only might lead to the conclusion that the increase in HYV was driving this phenomenon. However, by looking at a large sample of Indian districts over time and using control groups in rural and urban areas, I can distinguish the general direction of gender differences in rural areas and the performance associated with a deeper incidence of the Green Revolution across India.

3.3 Gender-specific Crops and Literacy

In the previous section I found that, on average, the introduction of Green Revolution seeds are not associated with lower measure of female status, once the all-rural India trend is controlled for. However, since the adoption of different crops of HYV can be associated with different intensities in female labour (both in absolute and relative terms to male labour), a breakdown of effects by crop can help depicting a more precise picture of the distributional effects of the new agricultural technology. In particular, as mentioned above, I would expect positive effects on literacy for boys and girls in places using more of the female-intensive crops (no conflict between the effect of returns to education and the intra-household allocation effect). However, in places where male-intensive crops becomes relatively more widespread, it is not clear which effect would prevail, in particular for girls.

To test this in the Indian context I make use of information on the two most important crops both in area cropped and HYV adoption, i.e rice and wheat[8]. To characterise the relative use of female labour in districts with different crops I make use of cross-sectional variation in female agricultural labour in districts with different

[8]Before the start of the Green Revolution, 20% of the cultivated area was cropped with rice and 11% with wheat, on average. Jowar, a variety of sorghum, followed with around 10% of total cropped area. In 1981, around 11% of the total cropped area was cultivated with HYV of wheat and another 10% with HYV of rice. The use of HYV Jowar was only 2% of total cropped area.

intensities of rice and wheat before the start of the Green Revolution, controlling for state fixed effects and district characteristics. Table 13 shows results for two different specifications using four different labour market indicators. In the first specification, I use the proportions of area cropped with rice and wheat. In the second, I construct a dummy for districts where the area cultivated with rice dominates the area used for wheat. In columns (1) and (2) I find that the proportion of women that do not work are lower the greater the intensity in rice. A greater proportion of wheat is associated with more women not working. Column (2) shows that districts that are rice intensive show around 9 percentage points more female labor. Columns (3) to (6) show that rice intensity draws more women to the agricultural sector, both as a proportion of total rural women and as a proportion of rural female workers. Wheat intensity shows the opposite effect. Less women work in rural areas and, if they do, they don't go to the agricultural sector. Note that these regressions control for characteristics such as the size of the manufacturing sector. Finally, columns (7) and (8) show that the ratio female - male in the agricultural sector also increases with rice intensity and decreases with wheat intensity.

Results in Table 13 suggest that, within states, rice was relatively more intensive in female labour. The question is whether improvements in literacy (overall and gender specific) are associated with the introduction a specific type of HYV seeds. If improvements in literacy come from increased returns to education associated with the technology, I should find that the introduction of HYV seeds improves literacy regardless of the gender-specificity of the crop. If improvements in literacy come from improvements in female relative income, as in Qian (2008), then I should find stronger effects in rice intensive areas or in areas where female labour was more prevalent before the start of the Green Revolution. Table 14 shows evidence of both phenomena being at work. Column (1) shows that the adoption of HYV seeds had a positive effect on rural literacy for the cohort of children aged between 5 and 14 irrespective of whether the HYV seed was rice or wheat. Furthermore, since the mean adoption was the same for both crops (around 10%) the magnitude of the effect is similar for both crops. A difference in difference analysis in column (2) shows that districts that were relatively rice intensive before the Green Revolution (and where women where more involved in the agricultural sector as shown by Table 13) have improved literacy less than districts that were relatively wheat intensive by 1971. Improvements in literacy were similar for both sets of districts in 1981.

Colums (3) and (5) show interesting results in terms of the heterogeneity of the crop effects by gender. Adoption of HYV of wheat seem to have a symmetric effect across genders. However, the effect on boys is slightly larger than the effect of HYV of rice on the same group and the effect on girls, slightly smaller. It follows that a greater adoption of HYV of rice would have a greater impact on girls' literacy than on boys'. Statistically, column (7) shows that the reduction in the gap in literacy between boys and girls is only associated with more adoption of high yield varieties of rice.

When comparing districts according to their pre-Green Revolution crop intensity, columns (4) and (6) show that the level effects on literacy for both boys and girls were greater in districts that were initially more intensive in wheat. However, column (8) shows that the gender gap was reduced only in areas that were initially intensive in rice after the start of the Green Revolution.

This section's results show that the returns to education effect was strong and that the Green Revolution was associated with improvements in the level of literacy across genders regardless of the gender-labour intensity of the crop. In particular, areas intensive in wheat seem to have gender-neutral improvements in literacy. However, the relative improvement of female literacy was associated with greater adoption of the new variety of the female-intensive crop (i.e. rice) and in areas where rice was initially more prevalent. That means that areas where female labour was relatively more important experienced greater improvements in female literacy. In other words, the relative impact of HYV adoption of a given crop on literacy was stronger for the gender whose labour is relatively more intensive in that crop, suggesting the presence of an effect derived from intra-household relative reallocation of resources and power.

3.4 Sex Imbalance

The problem of "missing women" in rural India, as described by Sen (1990), is observed in this paper's dataset, where the ratio of boys to girls aged 0 to 4 is above 1.02, on average. That means that changes in rural India that affect the status of women could also have an impact on the survival rate of girls relative to boys. Qian (2008), for example, showed that increases in the fraction of surviving girls in rural China followed improvements in female status. However, as in the case of female labour participation, raw numbers suggest a worsening in female status in all rural India, since the ratio of boys to girls increases over time (from 1.01 in 1961 to almost 1.03 in 1981). In this section I investigate whether a deepening of the Green Revolution can be associated with this trend or not, reproducing the same specifications used in the previous sections, this time using the ratio of boys to girls aged between 0 and 4 as the explained variable. I also check whether urban and older rural cohorts show a similar trend, as a placebo test.

Table 15 shows in columns (1) and (2) that year dummies are positive and significant, conveying the negative trend in the rate of surviving girls in rural India. However, an increase in HYV adoption is associated with a less negative evolution of the ratio of boys to girls. As in previous results, the Green Revolution only mitigates this trend, since only a value of HYV Adoption of 1 (i.e. a district using only HYV seeds) would counter the negative trend fully. Column (2) also shows that the mitigating effect of HYV adoption only shows up in 1981, suggesting that the survival rate of girls took longer to improve than other indicators of female status. Columns (3) and (4) perform the same regressions on the same cohort in urban areas and older cohorts, both cohorts supposedly unaffected by the introduction of the new agricultural technology. Column (3) shows that the worsening of the boys to girls ratio is only a rural phenomenon and that districts with more HYV adoption did not show a differential effect in urban areas. Column (4) shows that older cohorts in rural areas were on a steeper decline in the rate of girls' survival relative to boys' (as shown by the coefficients on the year dummies), but that there was no pre-existing differential trend for areas that subsequently showed a greater adoption of the new seeds. In both cases, the coefficients on HYV adoption were positive but not significantly different from 0.

4 Instrumental Variables Results

An additional concern with respect to literacy and female status results is that of endogeneity of HYV adoption. It might be the case that whenever younger people are more educated they take more control of farming decisions and, simultaneously through preferences (e.g. less risk averse) or information, they end up adopting new technologies. In that case, diffusion of HYV seeds could be a result and not a driver of improvements in literacy. Alternatively, unobserved changes in districts (e.g. cultural transformation) could explain why women are more active in income generating activities and why adoption of the new seeds diffuses more rapidly. In that case, more dynamic districts could go through both transformations independently.

To address this issue, I look for geographic characteristics that can explain HYV adoption without having a direct effect on outcomes of interest. For example, Rud (2009) and Foster et al. (2008) show that, since HYV seeds are irrigation intensive, a geological measure of aquifer thickness can explain their diffusion. Similarly, Evenson and McKinsey (1999) have produced a set of correlations that link geography, infrastructure and technology adoption in Green Revolution India. To identify which geographical variables can predict my measure of HYV adoption, I run a regression of HYV adoption on a set of dummies capturing the thickness of the aquifer (i.e. groundwater availability) ranging from very low (i.e. no aquifers) to high (i.e. aquifer thicker than 150 mts.); the relief of land surface (from low, i.e. flat, to high, i.e. elevated) and the topsoil depth (from very low, i.e. less than 50 cms thick, to high, i.e. greater than 300 cms thick). I control for time varying district characteristics and for state and year fixed effects. Table 16a shows that the best predictors of HYV adoption are the availability of groundwater (medium and high) and a flat land surface, relative to the respective categories omitted in the regression, namely very low groundwater availability and a very sloped terrain, respectively. There are no systematic differences in adoption of HYV seeds across different depths of topsoil.

As these characteristics are time invariant, in order to capture their time-varying effect on the diffusion of the new seeds I interact them with year dummies to predict HYV adoption in the first stage, as in Rud (2009). Column (1), Table 16b shows the instrumental variables first stage. Each of these geographic characteritics has a positive and increasing effect on HYV adoption. All interaction terms are positive and significant. Year dummies capture the increase in HYV adoption that is not explained by these characteristics. However, instruments look strong enough, since the F-test of the first stage is 23.44. In particular, districts with the thickest aquifers are those adopting more HYV seeds, as suggested by the magnitude of the coefficients on "Groundwater High".

The rest of Table 16b shows a set of reduced form regressions of the outcomes of interest where the endogeneity of HYV adoption was a concern. I also include outcomes that should not be affected by these geographic characteristics as a placebo test. Column (3) shows that rural literacy among rural children aged between 5 and 14 has increased in districts with more groundwater available. Not a single of the geographic characteristics seem to be associated with improvements in the literacy levels for the same cohort in urban areas or for older cohorts in rural areas, as shown in columns (3) and (4). That means that districts with greater groundwater availability only predict better literacy outcomes for rural cohorts in educational age after the start of the Green Revolution.

Columns (5) to (8) show results for female labour market outcomes. Across the board, coefficients on year dummies suggest that women participation in the labour market had been deteriorating over time (in absolute terms or relative to men's). However, this trend was less pronounced in districts with more groundwater availability or flatter terrain.

Table 17 presents the second stage regressions when adoption of HYV seeds is instrumented using year-geographic characteristics interactions. Columns (1) to (3) look at literacy, overall and by gender. Coefficients are positive and significant and magnitudes very similar to OLS regressions. This suggests that endogeneity was not a major issue in the literacy regressions or that, if present, opposing biases were cancelling out. Columns (4) to (7) show results for female labour outcomes and shows that HYV adoption has increase overall female participation in the labour market, and in the agricultural sector in particular, both in absolute terms and relative to males. Results are consistent with those found in OLS regressions in Table 9, even though coefficients are greater in magnitude, suggesting that the effect of the agricultural technology on female status suffered a negative bias. That could be explained by an unobserved time-varying phenomenon that is positively correlated with HYV adoption and negatively correlated with female labour, or viceversa. Overall, IV results alleviate concerns that the direction of causality was from literacy and female status to technology adoption.

5 Conclusions

The economic and social changes in rural areas associated with the introduction of high yield seeds in mid-1960s in India–known as the Green Revolution–still generate conflicting evidence and controversy. The overarching concern is that in the presence of strong heterogeneous effects, only the more prosperous or the more educated farmers benefited from the introduction of a new technology and many remained excluded. In particular, the prerequisites for successful adoption of the new technology, such as mechanization, irrigation and fertilizers, seemed to be very restrictive for a country with more than 70% rural population and high levels of poverty, landlessness and illiteracy.

This paper shows that, despite the possibility of heterogeneous effects within districts, the Green Revolution was associated with sizeable aggregate effects on rural literacy at the district level. Even though pre-HYV means show that adoption was more prevalent in districts with higher levels of rural literacy, the effects identified capture changes over and above the average increase in rural literacy experienced between census years. A district at the mean of the distribution of HYV intensity experienced an extra increase in rural literacy of around 1 percentage point in 1971 and more than 2 percentage points in 1981 with respect to pre-Green Revolution levels. Some identification concerns regarding the link between literacy and HYV adoption are alleviated by the lack of evidence that cohorts whose educational decisions should not have been affected by the level of HYV intensity experienced the same changes in literacy. I find no evidence of pre-Green Revolution trends in rural areas (as captured by older rural cohorts) or simultaneous district wide changes (as captured by the cohort of young urban individuals). Additionally, the identifying assumption–i.e. that these changes would not have happened by other means in the absence of the Green Revolution–is robust to the inclusion of time-varying district characteristics, state investment in primary schools and time trends.

The concerns that the female population might have been left out of the benefits of technological change are not supported by the data in terms of literacy rates and labour participation, on aggregate. Even though there is evidence that the average number of rural women excluded from labour markets increased after the Green Revolution, I find that this phenomenon was less prevalent, the greater the intensity in HYV. Similarly, the gender gap in rural literacy rates and the survival rates of girls increased over time since the introduction of HYV seeds, but the gap widened less in districts that adopted more HYV seeds. Without dismissing the possibility of heterogeneous effects among rural women, the aggregate data presents a picture that reconciles the conflicting anthropological evidence around the status of women after the start of the Green Revolution.

In many cases, the provision of infrastructure such as schools and teachers may be the leading constraint to education and its provision could spur increases in educational levels (see Duflo (2001), for example). But in other cases, low education may be the natural consequence of economic activities with low returns to education, such as traditional farming. The improvement in literacy for men and women and the reduction of the gender gap associated with the adoption of a new agricultural technology in rural areas, in particular when a relative female-intensive crop is used, suggest that improvements in human capital can be obtained in rural areas, if the returns to education increase through technological change.

References

[1] Arora, D.R., Agarwal, B.K. and Gupta, A.K.(1989), Green Revolution and Social Change in Punjab, in Green Revolution and Social Change, edited by Sharma and Dak,Ajanta Publications, Delhi.

[2] Bharadway, K. (1990), Irrigation in India: Alternative Perspectives, Indian Council of Social Science Research.

[3] Bertrand, M., Duflo, E. and Mullainathan, S. (2004), How Much Should We Trust Differences-in-Differences Estimates?, Quarterly Journal of Economics, 119(1), 249-75.

[4] Besley, T. and Burgess R. (1999), Land Reform, Poverty Reduction and Growth: Evidence from India, Quarterly Journal of Economics, 115(2), 389-430.

[5] Boserup, E.(1990), Economic Change and the Roles of Women, in Tinker I. (ed.) Persistent Inequalities: Women and World Development. Oxford: Oxford University Press.

[6] Burgess, R and Pande, R. (2005), Can Rural Banks Reduce Poverty? Evidence from the Indian Social Banking Experiment, American Economic Review, 95 (3), 780-795.

[7] Chakravarti, A.K. (1972), Green Revolution in India, Annals of the Association of American Geographers, 63 (3), 319-330.

[8] Dhanagare, D.N. (1989), The Green Revolution and Social Inequalities in Rural India, in Green Revolution and Social Change, edited by Sharma and Dak, Ajanta Publications, Delhi.

[9] Duflo, E.. (2001), Schooling and Labor Market Consequences of School Construction in Indonesia: Evidence from an Unusual Policy Experiment, American Economic Review ; 91(4), 795-813.

[10] Duflo, E. and Pande, R. (2007), Dams, Quarterly Journal of Economics, 122 (2), 601–646.

[11] Evenson, R.E. and McKinsey, J.W. (1999), Technology-climate interactions in the Green Revolution in India, Economic Growth Center Discussion Paper 805, Yale University.

[12] FAO (1997), Women and Food Security, FAO FOCUS on the Issues series.

[13] Foster, A and Rosenzweig, M. (1996), Technical Change and Human Capital Returns and Investments: Evidence from the Green Revolution„ American Economic Review 86(4), 931-953.

[14] Foster, A. and Rosenzweig, M. (2003a), Agricultural Productivity Growth, Rural Economic Diversity and Economic Reforms: India, 1970-2000, mimeo, Brown University.

[15] Foster, A. and Rosenzweig, M. (2003b), Agricultural Development, Industrialization and Rural Inequality, mimeo, Harvard University.

[16] Foster, A. and Rosenzweig, M. (2008), Inequality and the Sustainability of Agricultural Productivity Growth,: Groundwater and the Green Revolution in Rural India, mimeo.

[17] Gulati, A. and Narayanan, S. (2003), The Subsidy Syndrome in Indian Agriculture, Oxford University Press.

[18] Harriss J. (1989), Knowing About Rural Economic Change, in Bardhan P. Conversations Between Economists and Anthropologists: Methodological Issues in Measuring Economic Change in Rural India. Oxford: Oxford University Press.

[19] Jalan, J. and Ravallion, M. (2002), Geographic Poverty Traps? A micro model of consumption growth in rural China, Journal of Applied Econometrics, Vol. 17, pp 329-346.

[20] Kaur M. and Sharma, M.L. (1989), Impact of Green Revolution on Status of Rural Women, in Sharma & Dak (eds.), Green Revolution and Social Change, Ajanta Publications, Delhi.

[21] Kohli, D. S. and Singh, N. (1997), The Green Revolution in Punjab, India: The Economics of Technological Change, mimeo.

[22] Matsuyama, K. (1992), Agricultural Productivity, Comparative Advantage and Economic Growth, Journal of Economic Theory, Vol. 58, 317-334.

[23] Mazumdar and Sharma (1990), Sexual Division of Labour, in Tinker I. (ed). Persistent Inequalities: Women and World Development. Oxford: Oxford University Press.

[24] Mencher J. (1988), Women's Work and Poverty: Women's Contributions to Household Maintenance in Two Regions of South India, in Bruce J. and Dwyer D. (eds) A Home Divided: Women and Income Control in the Third World. Stanford: Stanford University Press.

[25] McGuirk, A and Mundlak, Y. (1991), Incentives and Constraints in the Transformation of Punjab Agriculture, International Food Policy Research Institute.

[26] Ministry of Agriculture and Irrigation, Government of India (1976), Report of the National Commission on Agriculture, Volume X.

[27] Planning Commission, India (1961), Third Five Year Plan.

[28] Qian, N. (2008), Missing Women and the Price of Tea in China: The Effect of Sex-Specific Earnings on Sex Imbalance, Quarterly Journal of Economics, Vol. 123 (3), 1251–1285.

[29] Ravallion, M. and Datt, G. (1999), When is Growth pro-poor? Evidence from the Diverse Experience of India's State, Policy Research Working Paper 2263, World Bank.

[30] Rosenzweig M. and Schultz T.P. (1982), Market opportunities, Genetic Endowments, and Intrafamily Resource Distribution: Child Survival in rural India, American Economic Review, Vol. 72(3), 803-815.

[31] Rud, J.P. (2009), Electricity Provision and Industrial Development: Evidence from India, mimeo, London School of Economics.

[32] Sen, A. (1990), More than 100 million women are missing, New York Review of Books, Vol. 37 (20),

[33] Sims, H. (1988), Political Regimes, Public Policy and Economic Development: Agricultural Performance and Rural Changes in Two Punjabs, Sage Publications, New Delhi.

[34] Sharma, M.L. and Dak, T.M (1989), Green Revolution and Social Change, Introduction, edited by Sharma and Dak,Ajanta Publications, Delhi.

[35] Sridhar, D. (2004), The Green Revolution in India: The Changing Nature of Work and its effect on child nutrition, University of Oxford, mimeo.

[36] Vanneman, Reeve and Douglas Barnes.(2000), Indian District Data, 1961-1991: machine-readable data file and codebook. College Park, Maryland: Center on Population, Gender, and Social Inequality.

[37] Wooldridge, J. (2002), Econometric Analysis of Cross Section and Panel Data, MIT Press.

[38] Wooldridge, J. (2003), Cluster-Sample Methods in Applied Econometrics, American Economic Review, XCIII, pp 133-138.

[39] World Bank (2000), India: Policies to Reduce Poverty and Accelerate Sustainable Development, World Bank.

[40] World Bank (2007), World Development Report 2008: Agriculture for Development, World Bank.

Appendix A: Data

The census data for years 1961-1981 comes from the Indian District Database prepared by R.Vanneman and D.Barnes (2000), Indian District Data, 1961-1991: data file and codebook, College Park, Maryland: Center on Population, Gender, and Social Inequality, available at www.inform.umd.edu/~districts/index.html. I merged the data with the India Agriculture and Climate Data set, prepared by Apurva Sanghi, K.S. Kavi Kumar and James W. McKinsey, Jr. for the World Bank (available at chd.ucla.edu/dev_data/datafiles/india_agric_climate.htm), accounting for changes in districts' boundaries that left me with 254 districts in 13 states (Andhra Pradesh, Bihar, Gujarat, Haryana, Karnataka, Maharashtra, Madhya Pradesh, Orissa, Punjab, Rajasthan, Tamil Nadu, Uttar Pradesh and West Bengal).

"HYV adoption" is the proportion of area cropped with HYV seeds and "HYV intensity" in the difference in difference tables is a dummy equal to 1 for all districts with HYV adoption above the median in a particular year. "Area Cropped with..." refers to the land cropped with the described seed as a proportion of total cropped area. "Rural literacy" for all age groups, gender and location is the proportion of literate people within that group. "Female population" and "Rural population" are the proportion of women and individuals living in rural areas in total population, respectively."Ratio Boys - Girls" is the number of boys over the number of girls in a given cohort and location. "Rural workers in manufactures" is the proportion of main workers in rural areas working in the manufacturing sector, and "Log Roads" is the log of roads per square kilometre. "Rural migration" is the proportion of population in a district that has come to rural areas from other districts. "Scheduled Caste" and "Scheduled Tribe" are the proportion of population that belongs to either group. I also include "Primary schools", which is the log of primary schools per capita at the state level, and I alternate it with the log of number of primary school teachers per capita at the state level and the log of state expenditure on primary schools per capita. Finally, "Rural women not workers" is the proportion of women that are not employed in any main or marginal activity. "Rural women in agriculture" is the number of women working in the agricultural sector taken either as a proportion of rural female, of rural female workers or as a ratio of the number of men working in agriculture. Geographic characteristics are all dummies: "Groundwater" is a geological measure of watertable thickness, divided in a set of 4 mutually exclusive dummies. It can be "very low", eaul to 1 if the district has no thick aquifers; "low", if the thickest aquifer is less than a 100 mts thick; "medium" if between 100 and 150 mts thick or "high" if greater than 150 mts. Similarly, "Slope" is characterised by 3 dummies, according to the relief of the terrain, ranging from very flat to highly sloped terrain (other details not provided in the original dataset). "Topsoil" refers to the depth of the upper layer of soil and is divided in 4 mutually exclusive dummies. It starts from "very low" (equal to 1 if topsoil less than 50 cms thick), then "low" (if between 50 and 100 cms), "medium" (between a 100 and 300 cms) and "high" (more than 300 cms).

Appendix B: Tables

Table 1: Summary Statistics

Variable	Year	Mean	Std. Dev.	Min	Max
HYV Adoption (%)	1961	0	0	0	0
	1971	0.102	0.111	0	0.678
	1981	0.255	0.160	0	0.767
Rural Literacy (%)	1961	0.196	0.075	0.044	0.547
	1971	0.248	0.094	0.062	0.651
	1981	0.308	0.108	0.085	0.699
Male Rural Literacy (%)	1961	0.310	0.099	0.070	0.637
	1971	0.366	0.111	0.096	0.720
	1981	0.439	0.118	0.132	0.752
Female Rural Literacy (%)	1961	0.077	0.059	0.009	0.456
	1971	0.122	0.085	0.014	0.580
	1981	0.168	0.109	0.019	0.645
Rural Women Not Workers (%)	1961	0.655	0.154	0.394	0.979
	1971	0.798	0.117	0.445	0.991
	1981	0.750	0.143	0.472	0.994

Table 2: Mean Rural Literacy by Cohort and Year (1961-1971)

Panel A:	Rural Literacy						
	Aged 5 to 14			Aged 15 to 24			Difference in difference across cohorts
	HYV intensity (in 1971)			HYV intensity			
	High	Low	Difference	High	Low	Difference	
	(1)	(2)	(3)	(4)	(5)	(6)	(7) = (3) - (6)
1971	0.299 (0.010)	0.241 (0.008)	0.058 (0.013)***	0.385 (0.011)	0.336 (0.011)	0.049 (0.016)***	0.009 (0.005)**
1961	0.239 (0.008)	0.204 (0.007)	0.035 (0.010)***	0.273 (0.008)	0.226 (0.008)	0.047 (0.012)***	-0.012 (0.003)***
Difference	0.060 (0.013)***	0.037 (0.004)***	0.023 (0.006)***	0.112 (0.014)***	0.110 (0.014)***	0.002 (0.006)	0.021 (0.004)***

Panel B:	Urban Literacy		
	Aged 5 to 14		
	HYV intensity (in 1971)		
	High	Low	Difference
	(1)	(2)	(3)
1971	0.550 (0.010)	0.548 (0.008)	0.002 (0.011)
1961	0.489 (0.008)	0.495 (0.007)	-0.005 (0.011)
Difference	0.061 (0.012)***	0.053 (0.010)***	0.007 (0.005)

Robust standard errors in parentheses, * significant at 10%; ** significant at 5%; *** significant at 1%. "HYV intensity" is the net proportion of area cropped with HYV seeds. "High" is for districts above the median and "low" for those below. "Rural Literacy" and "Urban literacy" are the proportion of literate people in rural and urban areas, respectively, for different age groups.

Table 3: Mean Rural Literacy by Cohort and Year (1961-1981)

Panel A:	Rural Literacy						
	Aged 5 to 24			Aged 25 to 59			
	HYV intensity (in 1981)			HYV intensity			Difference in difference across cohorts
	High	Low	Difference	High	Low	Difference	
	(1)	(2)	(3)	(4)	(5)	(6)	(7) = (3) - (6)
1981	0.403 (0.012)	0.336 (0.011)	0.066 (0.016)***	0.252 (0.008)	0.232 (0.007)	0.020 (0.011)*	0.046 (0.008)***
1961	0.250 (0.008)	0.213 (0.008)	0.037 (0.011)***	0.169 (0.005)	0.154 (0.006)	0.015 (0.008)*	0.022 (0.004)***
Difference	0.153 (0.014)***	0.123 (0.013)***	0.029 (0.008)***	0.084 (0.009)***	0.079 (0.009)***	0.005 (0.004)	0.024 (0.006)***

Panel B:	Urban Literacy		
	Aged 5 to 24		
	HYV intensity (in 1981)		
	High	Low	Difference
	(1)	(2)	(3)
1981	0.643 (0.010)	0.650 (0.008)	-0.007 (0.011)
1961	0.536 (0.008)	0.535 (0.007)	0.001 (0.010)
Difference	0.108 (0.012)***	0.115 (0.010)***	-0.008 (0.006)

Robust standard errors in parentheses, * significant at 10%; ** significant at 5%; *** significant at 1%. "HYV intensity" is the net proportion of area cropped with HYV seeds. "High" is for districts above the median and "low" for those below. "Rural Literacy" and "Urban literacy" are the proportion of literate people in rural and urban areas, respectively, for different age groups.

Table 4: Literacy and HYV Adoption (OLS)

	(1)	(2)	(3)	(4)	(5)	(6)
	Panel A: HYV Adoption and Year Effects					
	Rural Literacy					
	Aged 5 to 14		Aged 15 to 24		Aged 5 to 24	
HYV Adoption	0.121	0.093	0.069	0.063	0.105	0.084
	(0.029)***	(0.026)***	(0.026)***	(0.024)***	(0.028)***	(0.025)***
1971	0.035	0.061	0.104	0.109	0.059	0.077
	(0.004)***	(0.013)***	(0.005)***	(0.013)***	(0.004)***	(0.012)***
1981	0.091	0.125	0.144	0.149	0.110	0.134
	(0.009)***	(0.025)***	(0.008)***	(0.025)***	(0.008)***	(0.024)***
District Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
District Controls	No	Yes	No	Yes	No	Yes
Observations	762	762	762	762	762	762
R-squared	0.95	0.96	0.96	0.97	0.96	0.97
	Panel B: Controls					
Female Population (%)		1.49		1.52		1.56
		(0.55)***		(0.54)***		(0.52)***
Rural Population (%)		-0.02		-0.03		-0.02
		(0.12)		(0.12)		(0.12)
Log Population		-0.17		-0.15		-0.16
		(0.06)***		(0.06)***		(0.06)***
Rural workers in Manufacturing (%)		1.76		1.63		1.75
		(0.25)***		(0.23)***		(0.23)***
Log Roads		0.03		0.02		0.03
		(0.006)***		(0.008)**		(0.006)***

Robust standard errors in parentheses, clustered at the district level. * significant at 10%; ** significant at 5%; *** significant at 1%. "HYV intensity" is the net proportion of area cropped with HYV seeds at the district level. "Rural literacy" is the proportion of literate people in rural areas for different age groups. "Female Population" and "Rural Population" are the proportion of women and individuals living in rural areas in total population, respectively. "Rural Workers in Manufacturing" is the proportion of main workers in rural areas working in the manufacturing sector and "Log Roads" is the log of roads per square kilometre.

Table 5: Literacy and HYV Adoption in Untreated Cohorts

	(1)	(2)	(3)	(4)	(5)	(6)
	Rural Literacy		Urban Literacy			
	Aged 25 to 59		Aged 5 to 14		Aged 5 to 24	
HYV Adoption	0.007	-0.0003	0.027	0.011	-0.019	-0.028
	(0.012)	(0.012)	(0.021)	(0.021)	(0.018)	(0.021)
1971	0.031	0.035	0.054	0.062	0.074	0.081
	(0.002)***	(0.007)***	(0.004)***	(0.011)***	(0.003)***	(0.010)***
1981	0.079	0.084	0.095	0.100	0.117	0.122
	(0.004)***	(0.014)***	(0.007)***	(0.022)***	(0.005)***	(0.020)***
District Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
District Controls	No	Yes	No	Yes	No	Yes
Observations	762	762	762	762	762	762
R-squared	0.97	0.97	0.93	0.94	0.95	0.96

Robust standard errors in parentheses, clustered at the district level. * significant at 10%; ** significant at 5%; *** significant at 1%. "HYV intensity" is the net proportion of area cropped with HYV seeds at the district level. "Rural Literacy" and "Urban Literacy" are the proportion of literate people in rural and urban areas, respectively, for different age groups. District controls are Female Population, Rural Population, Rural Workers in manufacturing and Log Roads.

Table 6: Literacy and HYV Adoption (Year Effects)

		(1)	(2)	(3)	(4)	(5)	(6)
		Rural Literacy					
		Aged 5 to 14		Aged 15 to 24		Aged 5 to 24	
HYV * 1971		0.110	0.104	0.026	0.045	0.086	0.089
		(0.029)***	(0.029)***	(0.027)	(0.026)*	(0.026)***	(0.026)***
HYV * 1981		0.125	0.089	0.081	0.069	0.110	0.082
		(0.031)***	(0.027)***	(0.028)***	(0.025)***	(0.030)***	(0.026)***
1971		0.037	0.060	0.108	0.110	0.061	0.077
		(0.004)***	(0.013)***	(0.005)***	(0.013)***	(0.004)***	(0.012)***
1981		0.091	0.126	0.141	0.147	0.109	0.135
		(0.009)***	(0.025)***	(0.008)***	(0.026)***	(0.008)***	(0.024)***
Additional effect of mean HYV (in percentage points)	1971	1.1	1.1	0.3	0.5	0.9	0.9
	1981	3.2	2.3	2.1	1.8	2.8	2.1
District Fixed Effects		Yes	Yes	Yes	Yes	Yes	Yes
District Controls		No	Yes	No	Yes	No	Yes
Observations		762	762	762	762	762	762
R-squared		0.95	0.96	0.96	0.97	0.96	0.97

Robust standard errors in parentheses, clustered at the district level. * significant at 10%; ** significant at 5%; *** significant at 1%. "HYV" is the net proportion of area cropped with HYV seeds at the district level. "Rural Literacy" is the proportion of literate people in rural areas for different age groups. District controls are Female Population, Rural Population, Rural Workers in Manufacturing and Log Roads.

Table 7: Literacy and HYV Adoption (Year Effects, Untreated Cohorts)

	(1)	(2)	(3)	(4)	(5)	(6)
	Rural Literacy		Urban Literacy			
	Aged 25 to 59		Aged 5 to 14		Aged 5 to 24	
HYV * 1971	-0.018	-0.011	0.057	0.052	-0.004	-0.006
	(0.012)	(0.012)	(0.030)*	(0.030)*	(0.023)	(0.026)
HYV * 1981	0.014	0.003	0.019	-0.0003	-0.023	-0.035
	(0.014)	(0.013)	(0.021)	(0.022)	(0.019)	(0.021)*
1971	0.034	0.036	0.051	0.059	0.073	0.079
	(0.002)***	(0.007)***	(0.005)***	(0.011)***	(0.004)***	(0.010)***
1981	0.078	0.082	0.097	0.106	0.118	0.126
	(0.004)***	(0.014)***	(0.007)***	(0.022)***	(0.005)***	(0.019)***
District Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
District Controls	No	Yes	No	Yes	No	Yes
Observations	762	762	762	762	762	762
R-squared	0.95	0.96	0.93	0.94	0.95	0.96

Robust standard errors in parentheses, clustered at the district level. * significant at 10%; ** significant at 5%; *** significant at 1%. "HYV" is the net proportion of area cropped with HYV seeds at the district level. "Rural Literacy" is the proportion of literate people in rural areas for different age groups. District controls are Female Population, Rural Population, Rural Workers in Manufacturing and Log Roads.

Table 8: Literacy and HYV Adoption (Year Effects, Additional Controls)

	(1)	(2)	(3)	(4)	(5)
	Rural Literacy: Aged 5 to 14				
HYV * 1971	0.098	0.085	0.078	0.098	0.045
	(0.030)***	(0.029)***	(0.030)***	(0.029)***	(0.027)*
HYV * 1981	0.092	0.086	0.072	0.085	0.064
	(0.027)***	(0.027)***	(0.030)**	(0.029)***	(0.023)***
1971	0.041	0.053	0.068	0.056	
	(0.025)*	(0.015)***	(0.013)***	(0.014)***	
1981	0.074	0.124	0.135	0.128	
	(0.033)**	(0.027)***	(0.026)***	(0.027)***	
Additional Controls	Rural Migration * year	Scheduled Caste * year	Scheduled Tribe * year	Primary Schools (State)	State Time Trends + Schools (State)
District Fixed Effects	Yes	Yes	Yes	Yes	Yes
District Controls	Yes	Yes	Yes	Yes	Yes
Observations	762	762	762	762	762
R-squared	0.96	0.96	0.96	0.96	0.97

Robust standard errors in parentheses, clustered at the district level. * significant at 10%; ** significant at 5%; *** significant at 1%. "HYV" is the net proportion of area cropped with HYV seeds at the district level. "Rural Literacy" is the proportion of literate people in rural areas for different age groups. District controls are Female Population, Rural Population, Rural Workers in Manufacturing and Log Roads. "Rural Migration" is the proportion of population in a district that has come to rural areas from other districts. "Scheduled Caste" and "Scheduled Tribe" are the proportion of population that belong to either group. "Primary Schools" is the log of primary schools per capita at the state level.

Table 9: Rural Women and Labour Market Outcomes

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Rural Women Not Workers (%)		Rural Women in Agriculture (% of rural females)		Rural Women in Agriculture (% of rural female workers)		Ratio Female - Male in Agriculture	
HYV Adoption	-0.086		0.085		0.218		0.110	
	(0.030)***		(0.029)***		(0.070)***		(0.047)**	
HYV * 1971		-0.193		0.187		0.206		0.316
		(0.046)***		(0.045)***		(0.079)***		(0.071)***
HYV * 1981		-0.055		0.056		0.222		0.049
		(0.030)*		(0.030)*		(0.072)***		(0.048)
1971	0.099	0.107	-0.071	-0.078	0.040	0.041	-0.145	-0.160
	(0.014)***	(0.014)***	(0.014)***	(0.014)***	(0.023)*	(0.023)*	(0.023)***	(0.023)***
1981	0.021	0.006	0.001	0.015	0.013	0.011	0.025	0.051
	(0.028)	(0.029)	(0.028)	(0.029)	(0.042)	(0.043)	(0.045)	(0.046)
District Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
District Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	762	762	762	762	762	762	762	762
R-squared	0.94	0.94	0.93	0.93	0.87	0.87	0.94	0.94

Robust standard errors in parentheses, clustered at the district level. * significant at 10%; ** significant at 5%; *** significant at 1%. "HYV" is the net proportion of area cropped with HYV seeds at the district level. "Rural Women Not Workers" is the proportion of women that are not employed in any main or marginal activity, "Rural Women in Agriculture" is the number of women working in the agricultural sector taken either as a proportion of rural female or of rural female workers. "Ratio Female - Male in Agriculture" is the number of women in the district working in the agricultural sector divided by the number of men in the same sector. District controls are Female Population, Rural Population, Rural workers in Manufacturing and Log Roads.

Table 10: Literacy and HYV Adoption (Effects by Gender)

	(1)	(2)	(3)	(4)	(5)	(6)
	Rural Literacy				Urban Literacy	
	Aged 5 to 24		Aged 25 to 59		Aged 5 to 24	
	Male	Female	Male	Female	Male	Female
HYV Adoption	0.064	0.101	-0.018	0.014	-0.020	-0.036
	(0.023)***	(0.033)***	(0.016)	(0.013)	(0.020)	(0.022)
1971	0.062	0.090	0.035	0.041	0.052	0.113
	(0.01)***	(0.02)***	(0.01)***	(0.01)***	(0.01)***	(0.01)***
1981	0.118	0.147	0.096	0.079	0.084	0.166
	(0.022)***	(0.031)***	(0.016)***	(0.016)***	(0.02)***	(0.02)***
District Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
District Controls	Yes	Yes	Yes	Yes	Yes	Yes
Observations	762	762	762	762	762	762
R-squared	0.97	0.95	0.98	0.94	0.94	0.96

Robust standard errors in parentheses, clustered at the district level. * significant at 10%; ** significant at 5%; *** significant at 1%. "HYV" is the net proportion of area cropped with HYV seeds at the district level. "Rural Literacy" and "Urban Literacy" are the proportion of literate people in rural and urban areas, respectively, for different age groups and genders. District controls are Female Population, Rural Population, Rural Workers in Manufacturing and Log Roads.

Table 11: Literacy and HYV Adoption (Year Effects by Gender)

		(1)	(2)	(3)	(4)	(5)	(6)
		Rural Literacy				Urban Literacy	
		Aged 5 to 24		Aged 25 to 59		Aged 5 to 24	
		Male	Female	Male	Female	Male	Female
HYV * 1971		0.074	0.101	-0.022	-0.001	0.006	-0.016
		(0.021)***	(0.030)***	(0.016)	(0.011)	(0.026)	(0.027)
HYV * 1981		0.062	0.102	-0.017	0.018	-0.028	-0.041
		(0.023)***	(0.035)***	(0.017)	(0.015)	(0.021)	(0.022)*
1971		0.061	0.091	0.036	0.042	0.050	0.111
		(0.011)***	(0.016)***	(0.008)***	(0.008)***	(0.009)***	(0.011)***
1981		0.119	0.147	0.096	0.077	0.088	0.169
		(0.022)***	(0.032)***	(0.016)***	(0.017)***	(0.018)***	(0.023)***
Additional effect of mean HYV (in percentage points)	1971	0.8	1.0	-0.2	0.0	0.1	-0.2
	1981	1.6	2.6	-0.4	0.5	-0.7	-1.1
District Fixed Effects		Yes	Yes	Yes	Yes	Yes	Yes
District Controls		Yes	Yes	Yes	Yes	Yes	Yes
Observations		762	762	762	762	762	762
R-squared		0.97	0.95	0.94	0.98	0.94	0.94

Robust standard errors in parentheses, clustered at the district level. * significant at 10%; ** significant at 5%; *** significant at 1%. "HYV" is the net proportion of area cropped with HYV seeds at the district level. "Rural Literacy" and "Urban Literacy" are the proportion of literate people in rural and urban areas, respectively, for different age groups and genders. District controls are Female Population, Rural Population, Rural Workers in Manufacturing and Log Roads.

Table 12: Gender Gap in Rural and Urban Areas

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Gender Gap				Rural Gender Gap - Urban Gender Gap			
	Aged 5 to 24		Aged 25 to 59		Aged 5 to 24		Aged 25 to 59	
	Rural	Urban	Rural	Urban				
HYV Adoption					-0.074		-0.024	
					(0.024)***		(0.019)	
HYV * 1971	-0.042	0.027	-0.042	-0.039		-0.068		-0.003
	(0.023)*	(0.016)*	(0.012)***	(0.015)**		(0.025)***		(0.020)
HYV * 1981	-0.059	0.017	-0.052	-0.022		-0.076		-0.030
	(0.027)**	(0.015)	(0.017)***	(0.017)		(0.025)***		(0.020)
1971	0.011	-0.063	0.025	-0.018	0.074	0.073	0.045	0.043
	(0.005)**	(0.004)***	(0.003)***	(0.004)***	(0.0051)***	(0.005)***	(0.005)***	(0.005)***
1981	0.051	-0.079	0.074	-0.018	0.130	0.130	0.090	0.091
	(0.010)***	(0.006)***	(0.006)***	(0.007)**	(0.009)***	(0.009)***	(0.009)***	(0.009)***
District Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
District Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	762	762	762	762	762	762	762	762
R-squared	0.83	0.93	0.95	0.90	0.90	0.94	0.90	0.94

Robust standard errors in parentheses, clustered at the district level. * significant at 10%; ** significant at 5%; *** significant at 1%. "HYV" is the net proportion of area cropped with HYV seeds at the district level. "Gender Gap" is the difference between the proportion of male and female that are literate for different age groups and in rural and urban areas. District controls are Female Population, Rural Population, Rural Workers in Manufacturing and Log Roads.

Table 13: Rural Female Labour Markets by Crop (in 1961: pre-Green Revolution)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Rural Women Not Workers (%)		Rural Women in Agriculture				Ratio Female - Male in Agriculture	
			(% of rural females)		(% of rural female workers)			
Area Cropped with Rice (%)	-0.083		0.104		0.112		0.186	
	(0.040)**		(0.064)*		(0.054)**		(0.075)**	
Area Cropped with Wheat (%)	0.301		-0.516		-0.279		-0.610	
	(0.073)***		(0.115)***		(0.098)***		(0.135)***	
Dummy Rice		-0.089		0.139		0.128		0.166
		(0.017)***		(0.026)***		(0.022)***		(0.031)***
State Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
District Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	254	254	254	254	254	254	254	254
R-squared	0.73	0.73	0.73	0.72	0.70	0.72	0.76	0.75

All regressions are for 1961 only. Standard errors in parentheses. * significant at 10%; ** significant at 5%; *** significant at 1%. "Area Cropped with Rice(Wheat)" is the net proportion of area cropped with rice (wheat) seeds at the district level. The omitted category is the proportion of area cropped with other products. "Dummy Rice" is equal to 1 if the area cropped with rice is greater than the area cropped with wheat and 0 otherwise, measured in 1961. "Rural Women Not Workers" is the proportion of women that are not employed in any main or marginal activity, "Rural Women in Agriculture" is the number of women working in the agricultural sector taken either as a proportion of rural female or of rural female workers. "Ratio Female - Male in Agriculture" is the number of women in the district working in the agricultural sector divided by the number of men in the same sector. District controls are Female Population, Rural Population, Rural Workers in Manufacturing and Log Roads.

Table 14 : Rural Literacy by Crop (Cohort: Aged 5 to 14)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Rural Literacy - All		Rural Literacy - Male		Rural Literacy - Female		Rural Literacy - Gender Gap	
Area Cropped with HYV Rice (%)	0.099 (0.026)***		0.063 (0.027)**		0.141 (0.028)***		-0.078 (0.018)***	
Area Cropped with HYV Wheat (%)	0.107 (0.025)***		0.102 (0.021)***		0.105 (0.033)***		-0.003 (0.024)	
Dummy Rice * 1971		-0.023 (0.005)***		-0.028 (0.006)***		-0.015 (0.006)***		-0.013 (0.005)**
Dummy Rice * 1981		-0.004 (0.006)		-0.011 (0.006)*		-0.004 (0.008)		-0.016 (0.006)**
1971	0.024 (0.004)***	0.043 (0.004)***	0.019 (0.004)***	0.039 (0.004)***	0.046 (0.005)***	0.108 (0.013)***	-0.010 (0.004)***	-0.006 (0.004)
1981	0.056 (0.008)***	0.077 (0.009)***	0.066 (0.008)***	0.086 (0.009)***	0.065 (0.010)***	0.184 (0.025)***	0.019 (0.007)***	0.021 (0.006)***
District Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
District Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	762	762	762	762	762	762	762	762
R-squared	0.96	0.96	0.95	0.96	0.95	0.95	0.82	0.82

Robust standard errors in parentheses, clustered at the district level. * significant at 10%; ** significant at 5%; *** significant at 1%. "Rural Literacy" is the proportion of literate people in rural areas for different the age group of 5 to 14 years old, by genders. "Gender Gap" measures the difference between male and femal literacy. "Area Cropped with HYV Rice(Wheat)" is the net proportion of area cropped with HYV rice (wheat) seeds at the district level. "Dummy Rice" is equal to 1 if the area cropped with rice is greater than the area cropped with wheat and 0 otherwise, measured in 1961. District controls are Female Population, Rural Population, Rural Workers in Manufacturing and Log Roads.

Table 15 : Sex Ratio and HYV Adoption

	(1)	(2)	(3)	(4)
	Ratio Boys-Girls			
	Aged 0 to 4		Aged 0 to 4	Aged 20 to 24
	Rural	Rural	Urban	Rural
HYV Adoption	-0.020 (0.011)*		0.020 (0.022)	0.036 (0.024)
HYV * 1971		-0.004 (0.017)		
HYV * 1981		-0.025 (0.013)*		
1971	0.020 (0.007)***	0.019 (0.007)**	0.016 (0.013)	0.033 (0.012)***
1981	0.026 (0.013)**	0.028 (0.014)**	-0.003 (0.024)	0.067 (0.026)***
District Fixed Effects	Yes	Yes	Yes	Yes
District Controls	Yes	Yes	Yes	Yes
Observations	762	762	762	762
R-squared	0.74	0.74	0.48	0.86

Robust standard errors in parentheses, clustered at the district level. * significant at 10%; ** significant at 5%; *** significant at 1%. "HYV" is the net proportion of area cropped with HYV seeds at the district level. "Ratio Boys-Girls" is the number of boys divided by the number of girls for different age groups and rural/urban. District controls are Female Population, Rural Population, Rural Workers in Manufacturing and Log Roads.

Table 16 a : Geographic Characteristics and HYV Adoption

	(1)	(2)
	HYV Adoption	
Groundwater Low	0.013 (0.010)	0.006 (0.010)
Groundwater Medium	0.044 (0.014)***	0.046 (0.015)***
Groundwater High	0.059 (0.017)***	0.049 (0.018)***
Slope Low	0.020 (0.010)**	0.019 (0.010)**
Slope Medium	-0.013 (0.008)	-0.010 (0.008)
Topsoil Low	-0.002 (0.009)	0.001 (0.009)
Topsoil Medium	-0.015 (0.012)	-0.015 (0.011)
Topsoil High	0.019 (0.013)	0.025 (0.013)*
Fixed Effects	State and Year	
District Controls	No	Yes
Observations	762	762
R-squared	0.67	0.68

Robust standard errors in parentheses, clustered at the district level. * significant at 10%; ** significant at 5%; *** significant at 1%. "HYV Adoption" is the net proportion of area cropped with HYV seeds at the district level. Omitted categories are "Groundwater Very Low", "Slope High" and "Topsoil Very Low"

Table 16b : Geographic Characteristics and HYV Adoption, Literacy and Female Labour

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	HYV Adoption	Literacy			Rural Women			
		Rural	Urban	Rural	Not Workers	in Agriculture, as a % of		
		Aged 15 to 24		Aged 25 to 59		Rural Females	Rural Female Workers	Males in Agriculture
Groundwater Medium * 1971	0.089 (0.020)***	0.017 (0.009)*	-0.012 (0.008)	-0.006 (0.003)*	-0.053 (0.013)***	0.057 (0.017)***	0.007 (0.016)	0.089 (0.023)***
Groundwater Medium * 1981	0.127 (0.020)***	0.018 (0.011)*	-0.018 (0.008)**	0.001 (0.005)	-0.018 (0.013)	0.005 (0.015)	-0.013 (0.019)	0.007 (0.022)
Groundwater High * 1971	0.163 (0.016)***	0.054 (0.010)***	0.002 (0.008)	-0.007 (0.004)	-0.086 (0.020)***	0.103 (0.022)***	0.089 (0.033)***	0.144 (0.041)***
Groundwater High * 1981	0.234 (0.027)***	0.053 (0.014)***	-0.017 (0.011)	0.007 (0.008)	-0.059 (0.016)***	0.063 (0.012)***	0.133 (0.035)***	0.082 (0.029)***
Slope Low * 1971	0.039 (0.017)**	-0.013 (0.007)*	-0.001 (0.007)	-0.008 (0.003)***	-0.030 (0.011)***	0.044 (0.016)***	0.040 (0.012)***	0.056 (0.021)***
Slope Low * 1981	0.053 (0.017)***	-0.013 (0.009)	-0.009 (0.007)	-0.014 (0.005)***	-0.004 (0.011)	0.010 (0.012)	0.034 (0.013)***	0.002 (0.019)
1971	0.044 (0.011)***	0.032 (0.004)***	0.051 (0.005)***	0.032 (0.002)***	0.179 (0.006)***	-0.178 (0.009)***	0.040 (0.007)***	-0.286 (0.012)***
1981	0.162 (0.016)***	0.077 (0.008)***	0.087 (0.008)***	0.069 (0.004)***	0.114 (0.009)***	-0.097 (0.011)***	0.053 (0.011)***	-0.134 (0.020)***
District Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
District Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	762	762	762	762	762	762	762	762
R-squared	0.85	0.96	0.94	0.97	0.94	0.95	0.92	0.95

Robust standard errors in parentheses, clustered at the district level. * significant at 10%; ** significant at 5%; *** significant at 1%. "HYV Adoption" is the net proportion of area cropped with HYV seeds at the district level. "Rural Literacy" is the proportion of literate people in rural areas for different age groups. "Rural Women Not Workers" is the proportion of women that are not employed in any main or marginal activity, "Rural Women in Agriculture" is the number of women working in the agricultural sector taken either as a proportion of rural female, of rural female workers or of rural males working in agriculture. District controls are Female Population, Rural Population, Rural Workers in Manufacturing and Log Roads. First Stage F-Test Statistic is 10.41.

Table 17 : IV Regressions

Literacy and Female Labour when Instrumenting HYV Adoption using Geographic Characteristics

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Rural Literacy			Rural Women			
	Aged 15 to 24				in Agriculture, as a % of		
	Total	Male	Female	Not Workers	Rural Females	Rural Female Workers	Males in Agriculture
HYV Adoption	0.103	0.098	0.099	-0.251	0.276	0.385	0.385
	(0.043)**	(0.040)***	(0.056)*	(0.054)***	(0.067)***	(0.112)***	(0.105)***
District and Year Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
District Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	762	762	762	762	762	762	762
R-squared	0.96	0.95	0.95	0.94	0.94	0.90	0.95

Robust standard errors in parentheses, clustered at the district level. * significant at 10%; ** significant at 5%; *** significant at 1%. "HYV Adoption" is the net proportion of area cropped with HYV seeds at the district level. "Rural Literacy" is the proportion of literate people in rural areas for different age groups. "Rural Women Not Workers" is the proportion of women that are not employed in any main or marginal activity, "Rural Women in Agriculture" is the number of women working in the agricultural sector taken either as a proportion of rural female, of rural female workers or of rural males working in Agriculture. District controls are Female Population, Rural Population, Rural Workers in Manufacturing and Log Roads. First Stage F-Test Statistic is 23.44.