

NBER WORKING PAPER SERIES

SOCIAL STATUS, ECONOMIC DEVELOPMENT AND FEMALE LABOR FORCE
(NON) PARTICIPATION

Kaivan Munshi
Swapnil Singh

Working Paper 32946
<http://www.nber.org/papers/w32946>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
September 2024

Research support from Cambridge-INET and the Keynes Fund at the University of Cambridge, and the Agence Nationale de la Recherche (ANR) under the EUR Project ANR-17-EURE-0010 is gratefully acknowledged. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2024 by Kaivan Munshi and Swapnil Singh. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Social Status, Economic Development and Female Labor Force (Non) Participation
Kaivan Munshi and Swapnil Singh
NBER Working Paper No. 32946
September 2024
JEL No. J16, J20, O1

ABSTRACT

This research provides a status-based explanation for the persistent increase in female labor force non-participation (FLFNP) that often accompanies economic development. This explanation is based on the idea that households or ethnic groups in developing economies can signal their wealth, and thereby increase their social status, by withdrawing their women from the labor force. If the value of social status or the willingness to bear the signaling cost is increasing with economic development, then this would explain the persistent increase in FLFNP. To provide empirical support for this argument, we utilize two independent sources of exogenous variation—across Indian districts in the cross-section and within districts over time—to establish that status considerations determine rural FLFNP. Our status-based model, which is used to derive the preceding tests, is able to match the increase in rural Indian FLFNP that motivates our analysis. Counterfactual simulations of the estimated model indicate that conventional policy prescriptions, such as a reduction in the cost of female education, could *raise* FLFNP by increasing potential household incomes and, hence, the willingness to compete for social status. The steep increase in female education in recent decades could paradoxically have increased FLFNP in India even further.


Kaivan Munshi
Department of Economics
Yale University
87 Trumbull Street
New Haven, CT 06511
and NBER
kaivan.munshi@yale.edu

Swapnil Singh
CEFER
Totoriu g. 4
Bank of Lithuania
Vilnius, LT-01121
Lithuania
ssingh@lb.lt


1 Introduction

What is the relationship between female labor force participation and economic development? Goldin's (1994) canonical *U*-hypothesis posits that there is an initial decline, on account of an income effect, reduced employment opportunities outside agriculture, or gender norms. This is followed by an increase in female labor force participation at later stages of economic development, as female education rises with an accompanying increase in labor market returns. While there is empirical support for the *U*-hypothesis in some contexts (Olivetti, 2013), it runs counter to the development experience in many parts of the world (Klasen, 2019). Consider, for instance, the Indian economy, which has been growing since the 1950's after centuries of economic stagnation. As observed in Figure 1a, per capita GDP has increased at an approximately constant rate for many decades now, but rural female labor force non-participation, which we refer to as FLFNP henceforth, continues to rise, with no evidence of a reversal in this trend. To add to the puzzle, we see in Figure 1b that (higher secondary) education levels for males and females have converged over time. Female education is often accompanied by an increase in labor force participation, but this does not appear to be the case in India.

Figure 1: GDP per capita, FLFNP, and education (India)



(a) GDP per capita and FLFNP



(b) Secondary school enrollment

Source: NSS, ILO, World Bank

Our explanation for the preceding facts is based on a mechanism – social status – that is the subject of a long-standing literature in economics. In this literature, social status provides preferred access to non-market goods and services (Cole et al., 1992; Postlewaite, 1998). Social status is increasing in wealth, but wealth is not observed by the external agents who are responsible for allocating the goods and services. Conspicuous consumption is one way to signal wealth (Veblen, 1899). Alternatively, choices that are visibly income reducing, such as avoidance of remunerative occupations or abstaining from manual work or business can be used as signals (Weiss and Fershtman, 1998). This latter strategy is especially relevant for our analysis because households or ethnic groups in developing economies could potentially withdraw their women from the labor force as a way of

signaling their wealth. If the value of status or the willingness to bear the signaling cost is increasing with economic development, then this would explain the persistent increase in FLFNP that we observe in Figure 1a. Moreover, an increase in female education will increase households' *potential* income, which increases the competition for social status, as elucidated in the model described below. This explains why the increase in education that we observe does not result in a decline in FLFNP; indeed, it could increase FLFNP even further as we will see, reconciling Figures 1a and 1b.

While our ultimate objective is to examine the dynamic relationship between economic development and FLFNP in India, mediated by the status mechanism, we must first establish that there is indeed a link between FLFNP and social status. The challenge, which is also faced by previous studies that explore the status mechanism; e.g. Charles et al. (2009), Bursztyn et al. (2018) is that its benefits are not directly observed. Our empirical strategy, taking the lead from these studies, is to derive conditions under which the value of status or the willingness to bear the signaling cost is predicted to be relatively large, and then document that FLFNP is relatively high under precisely those conditions. Both sources of variation that we use to link FLFNP and status – in the cross-section and over time – only apply to rural populations and, hence, the analysis in this paper is restricted to those populations.

While the received literature tells us why social status is useful, it does not tell us where this will be the case. In general, social status is most beneficial when it provides preferred access to goods and services that are more valuable and are rationed. This is more likely to be the case in a large local economy: extending a metaphor from Postlewaite (1998), a front row seat provides greater utility in a large church than in a small church. The size of a local economy in the pre-industrial period would have been determined by agricultural productivity, which, in turn, can be measured by population density (Ashraf and Galor, 2011). We expect that this will also be true at early stages of economic development. Local amenities will be of higher quality in densely populated (more productive) rural areas, where there is greater aggregate output to fund them and a larger population. At the same time, there will be greater competition for access to these amenities. If external agents, such as the staff in local public facilities, use the status mechanism to allocate scarce resources, then this implies that the value of status is increasing in population density. It follows that the willingness to bear the signalling cost, measured by FLFNP, will also be increasing in population density (instrumented by agricultural productivity).

While a positive association between FLFNP and population density could potentially be generated by the status mechanism, other explanations are also available. The advantage of focussing on India is that it is not enough to signal your wealth to achieve high status in Hindu society. A central tenet of Indian sociology, going back to Srinivas (1956, 1967) is that status-seeking households must also make particular consumption choices – vegetarianism and teetotalism – that are associated with ritual purity and were traditionally adopted by the high castes. If the value of social status in rural India is increasing with population density, as we posit, then both rural FLFNP *and* these behaviors, with their associated non-pecuniary costs, should be increasing in that variable. Based on the status game that we describe below, this should be true for the low castes who are attempting to improve

their status and for the high castes who seek to maintain their social position.

In Section 2 of the paper, we use data from the population census, the India Human Development Survey (IHDS), and multiple rounds of the National Sample Survey (NSS) to provide empirical support for each step of the argument laid out above. In particular, we show that there is a positive association between population density and (i) the size of the local economy, (ii) the quality of local amenities and the competition for these amenities, and (iii) FLFNP, vegetarianism, and teetotalism, separately by caste. In addition, FLFNP, vegetarianism, and teetotalism are higher on average for the high castes, which implies that they have higher status in equilibrium. This is consistent with the well documented fact that higher castes receive preferential treatment in public and private facilities (Munshi, 2019).¹

While the results with vegetarianism and teetotalism provide support for our preferred status-based interpretation of the positive association between FLFNP and population density, we also consider alternative mechanisms in Section 2: (i) a reduced demand for female labor in more densely populated districts, (ii) a reduced supply of female labor in those districts for non-status reasons, and (iii) an income effect. If demand-side effects are driving our results, then female wages should be declining in population density, whereas we uncover precisely the opposite association. Among the supply-side constraints that have been proposed in the literature, it has been hypothesized that marriage and accompanying home production (child care) could be responsible for the withdrawal of women from the workforce in developing economies (Goldin, 1994; Afridi et al., 2018). However, we fail to uncover an association between population density and either marriage rates or fertility. Women may also be less likely to work if they have less education (Heath and Jayachandran, 2017). Here, again, we find that women residing in more densely populated districts actually have more years of schooling. This takes us to a final supply-side explanation for the positive association between FLFNP and population density, which is based on patriarchal norms that determine women’s status within their households and, by extension, their decision-making power and autonomy (Srinivas, 1977; Basu, 1992; Chakravarti, 1993).

The presumption in the patriarchal norms literature is that women would like to work for pay, but their low status, on account of the norms, prevents them from exercising their preferences. A desire for social status and low women’s status are thus both positively associated with FLFNP, but they differ in one important respect: social status signals provide information about (unobserved) wealth to the wider community, whereas the patriarchal norms are enforced within the household. Providing support for the social status mechanism, recent experimental evidence from (urban) India indicates that making women’s work externally visible has a substantial negative effect on their labor force participation (Jalota and Ho, 2024). Complementing this finding, we find that women’s decision-making and autonomy within the household, which is unobserved by the wider community, is uncorrelated with population density in rural India, using survey data from both the DHS and the

¹The model that we describe below does not predict how this differential treatment will vary with population density (this information is unavailable in any case). What it does tell us is that both low castes and high castes will incur greater signaling costs, measured by FLFNP, vegetarianism, and teetotalism, to improve their status (treatment) in more densely populated districts.

IHDS. More importantly, none of the mechanisms we have considered above, or an income effect for that matter, can explain why lower castes are more likely to be vegetarian and to abstain from alcohol in more densely populated districts. These behaviors could, in principle, vary for the high castes if their caste identity varies with population density. For the low castes, however, the only possible explanation is status-based emulation of the high castes (since these behaviors are not traditionally associated with their own identity).²

While the consumption behaviors that we associate with higher status are specific to Indian society, the positive association between FLFNP and population density should apply more generally. As expected, we document the same positive association in Asian countries, with data from the ILO and the DHS, in Section 2.³ To the best of our knowledge, the robust positive association between rural FLFNP and population density (agricultural productivity) that we have uncovered has not been previously documented in the literature. While we are unaware of a non-status explanation for the evidence provided thus far, any cross-sectional analysis of this sort has its limitations. In particular, there is always the possibility that unobserved spatial heterogeneity is driving the results. To provide independent support for the hypothesized link between FLFNP and status, we thus proceed to model the status game in Section 3, and then derive resulting implications for variation in FLFNP *within* Indian districts over time.

In Hindu society, particularly in rural areas, a household’s identity will be based on its caste and its status will be determined by its caste’s social position (Srinivas, 1967). Since the status game thus plays out at the caste level, we assume, for analytical convenience, that the local population consists of two (caste) groups, with all households in a group having the same wealth or income endowment. Households derive utility from the consumption of market goods and from a non-market good that is allocated through the status mechanism. Status is increasing in wealth, but household wealth is unobserved by the external agents who are allocating the non-market good, and thus must be signaled by a costly choice. Each household chooses its signal independently, but its share of the non-market good is determined by the aggregate contribution of all its group members (which determines the group’s status). In the equilibrium of this game, the average signaling cost in the local population, which we associate empirically with FLFNP is (i) increasing in the per capita value of status, (ii) increasing in the mean income endowment, and (iii) decreasing in the income endowment gap between the groups. As a corollary to these results, we show that they also apply to each group separately.

The first theoretical result that we derive, with respect to the value of status, formalizes the cross-sectional tests described above. The second result is new, but FLFNP could be increasing in income without a role for status. The third result, which implies that both lower castes and

²Atkin et al. (2021); Agte and Bernhardt (2023) also associate vegetarianism and teetotalism with caste identity, documenting that these norms are more likely to be followed by the upper castes when their identity is more salient. However, neither of these studies considers the possibility that the lower castes could adopt upper-caste conventions, as a way of emulating them when status is a consideration.

³In contrast, we do not observe any association between FLFNP and population density in sub-Saharan Africa. Our explanation for this regional difference, drawing on the literature in anthropology (Goody, 1971) and economics (Mayshar et al., 2022) is based on the fact that vertical stratification is largely absent in African society. Without stratification, the status mechanism cannot be used to allocate resources across groups, as in Asia.

higher castes will compete more vigorously as the income endowment gap narrows is most useful for ruling out alternative explanations, and also distinguishes our model from previous analyses that incorporate a role for status. In Bursztyn et al. (2018); Atkin et al. (2021); Macchi (2023); Dupas et al. (2024) there is no reference group. Observed signals could thus reveal absolute rather than relative income. Charles et al. (2009) incorporate relative income in their analysis, but individuals are trying to distinguish themselves from their own (racial) group and, hence, conspicuous consumption is *increasing* in the gap between their own income and the group’s mean income.⁴ The set up of our model and the key result with respect to the income gap between groups is actually more closely related to models of conflict that have been proposed in the literature; e.g. Esteban and Ray (2011); Mitra and Ray (2014), and this is not a coincidence. Social status and conflict are alternative (costly) mechanisms to allocate resources between groups, and when these groups are not vertically separated, as is the case for Hindus and Muslims in India or tribes in Sub-Saharan Africa, then the latter mechanism will need to be employed.

We test the implications of the model in Section 4 with data from multiple rounds of the NSS. While incomes will be derived from labor and land in a rural (agrarian) economy, we focus on wage income for the core tests of the model. In each district-time period, the mean *potential* income, which corresponds to the mean income endowment in the model, is computed as the weighted average of the mean wage in each caste-gender category. The weight for each category is based on the size of its working-age population, regardless of the occupational status of its members. The caste-gap in the income endowment is similarly constructed as the difference between the caste-specific potential incomes. The status mechanism or any unobserved factor that shifts female labor *supply* will also affect the equilibrium wage, which, as noted, is used to construct potential incomes. We account for this reverse causation, as well as for omitted variable bias and measurement error, by constructing statistical instruments for potential incomes that are based on rainfall in each district-time period. Rainfall in a rural economy will determine wages and, by extension, potential incomes through the *demand* for labor and, hence, our instruments plausibly satisfy the exclusion restriction.⁵ Our estimates indicate that FLFP is increasing in mean potential income and decreasing in the caste-gap in potential income, net of district and time period effects. These results are obtained separately for the low castes and the high castes, as implied by the model. We are unaware of any non-status explanation for why an exogenous narrowing of the income endowment between caste groups in a district should result in an increase in FLFP in both groups.⁶

Having established a link between FLFP and social status with two independent sources of exogenous variation, we complete the analysis in Section 5 by estimating the structural parameters of the model. For this analysis, we extend the analytical model developed in Section 3 by introducing

⁴In Genicot and Ray (2017) and Kim et al. (2024), parents similarly derive utility when their children’s income (education) exceeds that of their peers, which increases expenditures on education in equilibrium. However, these costly expenditures are not used to signal unobserved wealth, as conventionally assumed in the status literature.

⁵As discussed in Section 4, this instrument also allows us to relax an assumption in our model and in our construction of potential incomes, which is that individuals are homogeneous at the caste-gender level in a given district-time period.

⁶The model generates additional predictions for the magnitude of the coefficients on mean potential income and the caste-gap in potential income, by caste. We are able to verify these implications as well.

education choices and by allowing wages to be determined endogenously. We find that the model fits the data very well, with respect to FLFNP, education, and wages, across districts in each NSS round, and over time. The positive association between FLFNP and population density that we estimate in the cross-sectional analysis is assumed to arise because the value of status is increasing in the latter variable. Although this value cannot be observed directly, our parameter estimates indicate that it is increasing in population density in each NSS round, as assumed in the model. With regard to the observed increase in FLFNP over time, three factors could potentially generate this trend in our model: the value of status, mean potential income, and the caste-gap in potential income. Based on our parameter estimates, an increase in the first two factors over time is responsible for the increase in FLFNP. Economic development will increase the size of the local economy, which, in turn, will increase the value of status. It will also increase incomes. While the factors that increase FLFNP are thus a natural consequence of the development process, we are nevertheless interested in identifying policies that would ameliorate this inefficient signaling.

One conventional policy prescription would be to invest in female education. We evaluate this policy by exogenously reducing the cost of education and find that FLFNP actually *increases* substantially. While this result may be surprising at first glance, it is easily interpreted through the lens of our model: the decline in the cost of education and, for that matter, any scheme that offers a monetary incentive for women to work will increase their households' potential incomes. This will, in turn, increase the competition for status and its accompanying signaling costs. The steep increase in female education over time that we documented at the outset in India, very likely increased FLFNP even further. While the preceding discussion tells us that standard prescriptions to increase female labor force participation may not be effective, and even backfire, in economies where status considerations are relevant, our model does provide an alternative solution. The second counterfactual policy simulation that we consider reduces the non-pecuniary constraints to female labor force participation, by weakening patriarchal norms for instance. This effectively increases the cost of withdrawing women from the workforce, without changing potential incomes, and our simulations indicate that this strategy would result in a substantial *decline* in FLFNP.

2 Cross-Sectional Evidence

2.1 Labor Force Non-Participation Across Indian Districts

The cross-sectional test linking rural FLFNP to status that we propose is based on the following sequential argument: (i) The size of a local economy at early stages of economic development is determined by agricultural productivity, which, in turn, can be measured by population density. (ii) Amenities will be of higher quality in more densely populated rural areas, where there is greater aggregate output to fund them and a larger population. At the same time, there will be greater competition for access to these amenities. (iii) If the status mechanism is used to allocate amenities, then this implies that the value of social status, and accompanying investments in social status

including FLFNP, will be increasing in population density.

Population density and agricultural productivity: The source of exogenous variation in the cross-sectional test is agricultural productivity. Since this is a crop-specific statistic, we measure overall productivity by population density, based on the assumption that more productive areas can support a larger population in the pre-modern period and at early stages of economic development. In our analysis, population density at the district level is derived from the 1951 population census, which is just around the time the Indian economy was starting to develop and is as far back as we can go. While it may be reasonable to assume that population density at this early stage of development was largely determined by agricultural productivity, this variable could, in principle, have been affected by other factors such as historical famines and conflicts in the district. When we report associations with respect to population density in regression tables, we thus always instrument for population density with potential crop yields and when we present figures, the population density variable is always predicted population density. The FAO GAEZ database provides potential yields for 42 crops at different levels of technology and irrigation. Following Galor and Özak (2016) we use low technology-rain fed agriculture to measure the crop yields, so that population density is predicted by exogenous geo-climatic conditions alone.

The size of the economy and population density: The core cross-sectional test of the association between FLFNP and population density (instrumented by potential crop yields) will be implemented at different points in time over the 1987-2011 period. The implicit assumption is that fixed and exogenous agricultural productivity determines the size of the economy, the value of status and, hence, FLFNP at each subsequent point in time.


We verify the first part of the preceding argument by estimating the association between the size of the local economy in 2011 – the end point of our analysis – and population density. The size of a local economy will be determined by aggregate output. While Indian districts are divided into urban and rural populations, output statistics are only provided at the district level. We thus measure the size of the rural economy by the value of agricultural output in each district. Services and manufacturing will typically complement agricultural output at early stages of economic development in any case, so the focus on agriculture is not necessarily a limitation. As seen in Appendix Figure B1a, the value of agricultural output in 2011 is increasing in 1951 population density (predicted by potential crop yields).

The quality of local amenities and population density: The next step in the argument is to verify that amenities are of higher quality in more densely populated districts. We verify that this is indeed the case with an illustrative example from the public health system. In the Indian rural health system, Primary Health Centers (PHC's) serve as the first point of contact with the population, followed by Community Health Centers (CHC's) and rural (sub-district) hospitals at the next level. Cases that cannot be handled within this system are referred to the district hospital.⁷ In

⁷Sub-centers at the very bottom of the hierarchy provide the most basic services, but these rudimentary facilities are typically not even staffed by doctors.

principle, each type of facility should provide the same level of service in all districts. In practice, we expect that more densely populated districts, which generate more resources, will be better served. The village directory of the 2011 population census provides information on the health facilities in each village, which can be aggregated up to the district level. We measure the size of a facility by the number of doctors in place and, as can be seen in Figure 2a, size is increasing in population density for each type of facility.⁸ While we would expect larger facilities to provide a wider range of services, this information is not available in the census. However, the 2011 round of the IHDS did collect information on both size (the number of doctors) as well as the services that were provided by all health facilities in the Primary Sampling Units (PSU's) that it covered. Focussing on PHC's, CHC's, and rural hospitals, we see in Appendix Table B1 that the number of procedures and tests, as well as the range of equipment, is increasing in size for each type of facility.

Figure 2: Quality and demand for medical facilities (rural India)



Source: 2011 population census, Village Directory (Asher et al., 2021) and 1951 population census. Population density in 1951, measured in logs, is predicted by FAO GAEZ potential crop yields. State fixed effects are partialled out prior to nonparametric estimation using the Robinson (1988) procedure. All variables are standardized by subtracting the mean and dividing by the standard deviation.

Bringing the preceding results together, individuals residing in more densely populated districts have access to a wider range of services for each type of medical facility. Since all facilities provide basic services, this is effectively saying that these individuals have access to more advanced services. At the same time, they face greater competition for access to local health facilities. As documented in Appendix Figure B1a, more densely populated districts (not surprisingly) had larger rural populations in 2011. While the rural health facilities are supposed to cover a fixed population in all districts and, hence, should be proportionately more numerous in densely populated districts, we see in Figure 2b that the number of facilities per capita is decreasing in population density in practice.⁹

⁸All variables in Figure 2 are standardized by subtracting the mean and dividing by the standard deviation. This allows us to plot the different types of facilities, which have very different levels, on the same figure.

⁹This is also true if we replace the number of facilities by the number of doctors. The number of doctors per capita

If the staff in the rural health facilities use the status mechanism to allocate scarce resources across patients, then the preceding results imply that the value of status is increasing in population density. It follows that the willingness to bear the signaling cost and, hence, FLFNP will also be increasing in population density. This argument is evidently not restricted to the health sector and will apply more generally.

While a positive association between FLFNP and population density could be generated by the status mechanism, as described above, other explanations are available. One reason why we focus on India in our analysis is because Hindu society, with its hierarchy of castes, is especially amenable to an analysis of status. The Hindu population is vertically stratified into broad caste categories or *varnas*, within which are numerous endogamous castes or *jatis*. Caste networks serving different economic roles have historically been organized and continue to be organized at the level of the *jati* (Munshi, 2019). For an analysis of status, however, it is the *varnas* that are relevant. All of the surveys that we use for the analysis in this paper indicate whether a household is Scheduled Caste (SC), Scheduled Tribe (ST) or unclassified. The SC and ST households lie at the bottom of the Hindu social hierarchy, and we will thus treat these households as *low* status, while all other Hindu households are treated as *high* status. While this ranking may be immutable, the relative social position of these groups could vary across space and over time. Srinivas (1956, 1967) describes a process of “Sanskritization” with economic development in which the low castes attempt to raise their social position by emulating behaviors traditionally associated with the high castes.¹⁰ Withdrawal of women from the labor force is one such behavior. Other behaviors are linked to ritual purity and, among them, Srinivas emphasizes vegetarianism and teetotalism. If the value of social status is increasing with population density, as we posit, then both rural FLFNP *and* these consumption choices should be increasing in that variable. This should be true for the upwardly mobile low castes and for the high castes who are attempting to maintain their social position, as modelled formally in Section 3.¹¹

Labor force non-participation and population density: We use the NSS Employment and Unemployment surveys for the analysis of labor force participation and for the supporting analyses of education and wages that follow. These surveys include repeated cross-sections of households over the 1987-2011 period, selected through stratified random sampling, that are representative of the country’s population in each round. The labor force participation statistics are derived from the usual activity status of all working-age adults in each sampled rural household (see Appendix A for details of variable construction). Individual responses are aggregated up to the district level in each survey round, by caste and gender where relevant, to construct the statistics that we use for the analysis.


is decreasing in population density, separately for each type of facility (see Appendix Figure B1b).

¹⁰The Scheduled Tribes are not Hindus *per se*, but Srinivas (1956, 1967) documents that these historically marginalized groups also attempt to improve their social position through Sanskritization.

¹¹The status game, as we model it in Section 3, is more appropriately played between *jatis* at the level of the village. The NSS does not provide village identifiers or *jati* information. Our analysis thus aggregates decisions from many underlying status games.

Figure 3a reports the nonparametric association between rural FLFNP and population density in the earliest available (1987) and last available (2011) NSS round.¹² FLFNP is increasing in population density, measured in 1951, across Indian districts in each round. The fraction of working age (18-65 year old) women who are withdrawn from the labor force ranges from 0.45 to 0.8 in 1987. While this enormous cross-sectional variation does decline over time, reflected in the flatter slope in 2011, notice that there is an overall increase in FLFNP from 1987 to 2011 (at all levels of population density). Labor force participation for the men (Figure 3b), in contrast, does not vary with population density.

Figure 3: Rural labor force non-participation (Indian districts, NSS)



Source: NSS and 1951 population census

Population density in 1951, measured in logs, is predicted by FAO GAEZ potential crop yields.

State fixed effects are partialled out prior to nonparametric estimation using the Robinson (1988) procedure.

There are eight NSS rounds over the 1987-2011 period. Table 1 pools all these rounds to estimate the association between labor force participation and population density, as well as the change in this association over time. Column 1 reports the association between FLFNP and population density.¹³ Matching Figure 3a, the population density coefficient is positive and significant, while the interaction with the time trend is negative and significant. Columns 2-3 replace FLFNP, measured across all rural households in each district-time period, with the corresponding statistics for high castes and low castes, respectively. The same pattern of coefficients is obtained, in contrast with the men, where all coefficients are close to zero in Columns 4-6. While FLFNP is increasing with population density for both caste groups, notice that it is higher on average for the high castes. This difference in levels implies that the high castes have higher status than the low castes in equilibrium.

We verify the robustness of the core results that we have presented in Table 1 in the following

¹²All the analyses with district-level data in this section of the paper control for state effects. These fixed effects are partialled out nonparametrically using the Robinson (1988) procedure, as described in Appendix B.

¹³Indian districts will often divide over time and we take account of this by measuring outcomes at the level of contemporaneous administrative boundaries in the analysis. However, standard errors are clustered at the level of the original 1981 boundaries and population densities are set at their 1951 levels, as discussed in Appendix A.

Table 1: Rural labor force non-participation (Indian districts, NSS)

Dep. variable	rural labor force non-participation					
	female			male		
Caste group	all (1)	high (2)	low (3)	all (4)	high (5)	low (6)
Population density	0.124*** (0.029)	0.132*** (0.031)	0.097*** (0.027)	0.001 (0.003)	0.000 (0.005)	-0.002 (0.006)
Population density $\times$ time trend	-0.003*** (0.001)	-0.004*** (0.001)	-0.002** (0.001)	0.000** (0.000)	0.000 (0.000)	0.000 (0.000)
Kleibergen-Paap F-statistic	22.61	28.16	22.35	22.61	28.20	22.21
Dep. var. mean	0.658	0.692	0.595	0.085	0.091	0.073
Observations	3418	3401	3368	3420	3404	3370

Source: NSS (“thick” and “thin” rounds) and 1951 population census

Population density in 1951, measured in logs, is instrumented using FAO GAEZ potential crop yields.

State and NSS round fixed effects are included in the estimating equation.

Standard errors are clustered at the level of 1981 district boundaries.

ways in Appendix B: (i) Afridi et al. (2018) document that FLFNP rates are especially high in the 25-65 age range, when virtually all women are married. We thus restrict the sample to the 25-65 age range in Table B2. (ii) Previous analyses utilizing NSS data in the economics literature have typically restricted attention to the “thick” rounds, conducted in 1987-1988, 1999-2000, 2004-2005, 2009-2010 and 2011-2012; e.g. Mitra and Ray (2014); Afridi et al. (2018); Atkin et al. (2021). We follow these studies and use the five “thick” rounds for the analysis of consumption that follows. For the analysis based on the Employment and Unemployment surveys, however, we also utilize data from three additional “thin” rounds, conducted in 2004, 2005-2006, and 2007-2008. This gives us more variation over time within districts, which is needed to test the model in Section 4. As a robustness check, we only include “thick” rounds in Table B3. (iii) Muslim, Christian and Sikh societies in India are also stratified (Ahmad, 1967; Luke and Munshi, 2011; Judge, 2002). We thus expect the status game to play out within these other religious groups as well, resulting in a positive association between FLFNP and population density, and this is indeed what we observe in Table B4.

Vegetarianism, teetotalism and population density: We use the NSS Household Consumer Expenditure surveys for the analysis of vegetarianism and teetotalism. Table 2 replaces FLFNP with vegetarianism and teetotalism as the dependent variables when estimating the association with population density (see Appendix A for a detailed description of the construction of these variables). Providing additional support for the status mechanism, the pattern of coefficients and the mean of the dependent variable across caste groups with these complementary outcomes matches what we obtained with FLFNP as the dependent variable. Variation in vegetarianism and teetotalism across districts could, in principle, be driven by standard determinants of consumption demand; i.e. income

Table 2: Rural vegetarianism and teetotalism (Indian districts, NSS)

Dep. variable	vegetarianism			teetotalism		
	all (1)	high (2)	low (3)	all (4)	high (5)	low (6)
Population density	0.043** (0.017)	0.042** (0.017)	0.049** (0.021)	0.087*** (0.020)	0.061*** (0.016)	0.084*** (0.023)
Population density $\times$ time trend	-0.002*** (0.000)	-0.002*** (0.000)	-0.002** (0.001)	-0.002*** (0.001)	-0.001** (0.000)	-0.002** (0.001)
Kleibergen-Paap F-statistic	23.59	22.67	15.53	23.59	22.67	15.53
Dep. var. mean	0.608	0.640	0.540	0.848	0.887	0.782
Observations	2083	2078	2068	2083	2078	2068

Source: NSS (“thick” rounds) and 1951 population census

Population density in 1951, measured in logs, is instrumented using FAO GAEZ potential crop yields.

State and NSS round fixed effects are included in the estimating equation.

Standard errors are clustered at the level of 1981 district boundaries.


and prices. Non-vegetarian foods are relatively expensive and our results could thus be obtained if household incomes (expenditures) were declining in population density. Alternatively, supply-side effects could result in higher prices for non-vegetarian food items and alcohol in more densely populated districts. As seen in Appendix Table B5, these alternative explanations do not appear to be relevant. Total expenditures and food expenditures per household are increasing in population density. Moreover, the relevant prices are (weakly) decreasing in population density. This last result is indicative of a reduced demand for these products, in line with the status mechanism.

Notice that the coefficient on the population density variable is positive and significant, while the coefficient on the population density-time trend interaction is negative and significant, without exception, in Table 1 with FLFNP as the outcome and in Table 2 with vegetarianism and teetotalism as outcomes. This consistency indicates that these outcomes are linked. We provide direct support for the preceding claim in Figure 4 by reporting the correlation between FLFNP and the complementary consumption behaviors, across districts and over NSS rounds. The binned scatter plots reported in the figure indicate that these correlations are indeed strongly positive, both for low castes and high castes, at all levels of FLFNP, to complete the cross-sectional tests of the status mechanism with Indian data.


2.2 Alternative Explanations

The discussion in this section considers alternative (non-status) factors that have been proposed in the literature as determinants of female labor force participation in developing countries, showing that they cannot explain the associations with population density that we have estimated.

Figure 4: Rural vegetarianism, teetotalism, and female labor force non-participation (Indian districts, NSS)



(a) Vegetarianism



(b) Teetotalism

1. Household income effects: With economic development, there will be an increase in household income. Female leisure or, equivalently, FLFNP could then rise on account of this income effect (Goldin, 1994). We saw above that household expenditures were increasing with population density and thus the positive association we have uncovered could be due to an income effect. However, this mechanism would not explain the positive association between population density and both vegetarianism and teetotalism that we have estimated. Moreover, an income effect would not generate the additional implications of our status model with regard to the variation in FLFNP within districts over time, that we discuss and verify in Sections 3 and 4.

2. Demand for female labor: The demand for female labor in agriculture will depend on geo-climatic conditions, as documented for India by Carranza (2014). The demand for female labor in the industrial sector will also vary at early stages of economic development (Goldin, 1994). If the overall demand is decreasing or less remunerative occupations are available in more densely populated districts, then this would explain why women residing in these districts are less likely to work, with an accompanying decline in the equilibrium wage. In contrast, if women are less likely to work due to a supply-side constraint associated with the status mechanism, then female wages should be *increasing* in population density. The NSS reports wages for women who work for pay (see Appendix A for details) and we see in Table 3, Column 1 that there is a positive and significant association between wages and population density. This positive association is also obtained separately for high caste and low caste women in Columns 2-3.

3. Female labor supply: We posit that female labor supply in more densely populated districts is constrained due to the status mechanism. However, there could be other constraints on labor supply. As noted, increases in female education with economic development have been seen to raise labor supply. If more densely populated districts have lower female education, then this could explain

Table 3: Rural female wages and education (Indian districts, NSS)

Dep. variable	mean log wage			mean log years of education		
	all	high	low	all	high	low
Caste group	(1)	(2)	(3)	(4)	(5)	(6)
Population density	0.123** (0.062)	0.089 (0.083)	0.140** (0.065)	0.463*** (0.078)	0.375*** (0.076)	0.358*** (0.107)
Population density $\times$ time trend	-0.003 (0.003)	-0.001 (0.003)	-0.005* (0.003)	-0.012*** (0.003)	-0.011*** (0.003)	-0.002 (0.004)
Kleibergen-Paap F-statistic	14.86	10.37	8.34	21.50	20.10	6.50
Dep. var. mean	2.303	2.380	2.202	0.960	1.146	0.428
Observations	3206	2893	2908	3408	3381	3109

Source: NSS (“thick” and “thin” rounds) and 1951 population census
Population density in 1951, measured in logs, is instrumented using FAO GAEZ potential crop yields.
State and NSS round fixed effects are included in the estimating equation.
Standard errors are clustered at the level of 1981 district boundaries.

the positive association between FLFNP and population density that we have uncovered. As seen in Table 3, Columns 4-6, however, female education (see Appendix A for details) is *increasing* in population density. This result is obtained for high castes and low castes.

Apart from human capital, demographic characteristics can also affect female labor supply. With economic development, the returns to home production (child rearing) increase, with a commensurate increase in FLFNP (Goldin, 1994; Afridi et al., 2018). If marriage rates or fertility rates are increasing in population density, then the observed positive association with FLFNP could be obtained without a role for status. While the NSS provides information on each adult and child in the household, it does not link mothers to their children.¹⁴ We thus turn to the DHS, which provides information on marriage and fertility for a nationally representative sample of women. We begin in Table 4, Column 1 by verifying that there is a positive and significant association between female unemployment and population density, measured at the district level, with the DHS data.¹⁵ However, there is no association between population density and either marriage rates or fertility (measured by the number of surviving children or the number of children ever born) in Columns 2-4.

4. Patriarchal norms: While the preceding discussion has focussed on economic and demographic factors, traditional patriarchal norms have also been seen to determine female labor force participation in India. These norms determine a woman’s status within her household, which, in turn, determines her decision-making power and autonomy (Srinivas, 1977; Basu, 1992; Chakravarti,

¹⁴The NSS household roster reports the relationship between the head and each member, but this does not link mothers to their children in joint families, which are common in India.

¹⁵The DHS collects information on employment rather than labor force participation (see Appendix A) but these variables are highly correlated in practice. Restricted-use DHS data, which we utilize for the analysis, provide geo-codes for each survey cluster, which can be mapped to the district in which it is located.

Table 4: Rural demographic characteristics and patriarchal norms (Indian districts, DHS)

Dep. variable	status signal	demographic characteristics			patriarchal norms		
	female unemployment	marriage rate	children ever born	children alive	health decisions	expenditure decisions	can visit relatives
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Population density	0.049** (0.022)	-0.009 (0.008)	-0.039 (0.031)	-0.042 (0.028)	0.009 (0.016)	0.002 (0.018)	0.008 (0.016)
Kleibergen-Paap F-statistic	7.08	4.55	4.69	4.69	7.08	7.08	7.08
Dep. var. mean	0.631	0.782	1.186	1.092	0.738	0.719	0.729
Observations	512	598	590	590	512	512	512

Source: 2015 DHS and 1951 population census

Marriage rates and fertility rates, measured by the number of children ever born and children alive, are measured in logs at the district level in Columns 2-4.

Patriarchal norms are measured at the district level by the fraction of women who have a say with regard to household decisions about health and expenditures, and who can visit their relatives without permission, in Columns 5-7.

Population density in 1951, measured in logs, is instrumented by FAO GAEZ potential crop yields.

State fixed effects are included in the estimating equation.

1993). The presumption in the patriarchal norms literature is that women would like to work for pay, but their low status on account of the norms, keeps them at home. Spatial variation in women's status could then explain the positive association between FLFP and population density. High caste women traditionally had low status within their households (Srinivas, 1977; Chakravarti, 1993). This would explain the additional observation that high caste women are less likely to work.

Visible indicators of low female status, such as withdrawal of women from the labor force and their seclusion in the home are also associated with high social status. One way to disentangle these potentially coexisting mechanisms is to examine decision-making and autonomy within the household, which are unobserved by the wider community (and hence unaffected by social status concerns). The DHS elicits information from female respondents about their decision-making power. As seen in Table 4, Columns 5-7, the fraction of women who report they have a say with regard to household decisions about health and expenditures, and who do not need permission to visit their relatives, is independent of population density. Appendix Table B6 reports estimates with measures of autonomy obtained from the India Human Development Survey (IHDS) where we see, once again, that there is no association with population density.¹⁶ There is no evidence that women's status is declining with population density, although we note that this evidence is based on a limited set of outcomes and the crop suitability instruments have less statistical power in Table 4.

Taking a different approach, our results on vegetarianism and teetotalism, particularly the positive association between these variables and population density for the upwardly mobile lower castes,


¹⁶The IHDS is a nationally representative survey of households that was conducted in 2007 and 2011. Data from the second round can be used to construct measures of FLFP, vegetarianism and teetotalism. Matching the core NSS results, we see in Appendix Table B7 that each of these variables is positively associated with population density. The IHDS also collects information on whether the adult women in the household are "veiled." This is a visible indicator of low female status *and* high social status and, as expected, it is positively associated with population density.

provide additional and independent support for the social status mechanism. Agte and Bernhardt (2023) exploit a different source of exogenous cross-sectional variation to document that upper castes are less likely to make choices that are traditionally associated with their high status – FLFNP, vegetarianism, teetotalism – when their incomes are relatively low.¹⁷ However, this could simply reflect a weakening of their caste identity, as in Atkin et al. (2021). To uncover the social status mechanism, an ethnic group must be seen to emulate the customary behaviors of a *higher* status group, when the value of status is high, as we document in our data.

2.3 Labor Force Non-Participation Across Regions

The positive association between FLFNP and population density (agricultural productivity) is not specific to India and should be observed in other developing economies. We thus proceed to examine the association between rural FLFNP and population density across countries, separately in South and South East Asia and in Sub-Saharan Africa in Figure 5a. Rural labor force participation rates are obtained at the country level in 2005 from the ILO UN STATS database and population densities, derived from the NASA SEDAC database, are measured in 2000 (see Appendix A for details). As above, we only use that part of the variation in population density that can be explained by exogenous crop suitability, obtained from the FAO GAEZ database, in our analysis.

Figure 5: Rural labor force non-participation across regions (country data, ILO)



Source: ILO UN STATS and NASA SEDAC


Population density in 2000, measured in logs, is predicted by FAO GAEZ potential crop yields.

As observed in Figure 5a, FLFNP is increasing steeply with population density across Asian countries. However, this relationship is not observed across African countries, where the association

¹⁷In our model, an exogenous narrowing of the income endowment gap increases FLFNP in both caste groups. While this would appear to be at odds with Agte and Bernhardt’s findings, we note that their analysis is situated in a very unusual setting in which Scheduled Tribes have substantially *higher* income than upper castes. Our model, in which upper castes always have higher incomes than lower castes, does not apply to such a setting.

is (if anything) mildly negative. Based on the figure, the well documented difference in FLFNP between these regions can be almost entirely explained by the positive association with respect to population density in Asia but not Africa. For the men, in contrast, these inter-regional differences are absent in Figure 5b and there is no association between labor force non-participation and population density in Asia or Africa.

Figure 6: Rural labor force non-participation across regions (district data, DHS)



Source: DHS and NASA SEDAC

Population density in 2000, measured in logs, is predicted by FAO GAEZ potential crop yields.

First administrative unit (state) fixed effects and survey year effects are partialled out prior to nonparametric estimation using the Robinson (1988) procedure.

We subject the preceding facts to closer scrutiny in Figure 6 and the accompanying Appendix Table B8 (which provides regression coefficients with standard errors) by using DHS data, which are available for eight Asian countries and 29 African countries at different points in time (Appendix Table A1 lists these countries and the survey years). Rural employment rates, which are closely related to labor force participation rates, can be constructed at the district (second administrative unit) level with these data. Although there is now much greater overlap in population densities across regions, the same patterns are observed: (i) there is a positive and statistically significant association between female unemployment and population density across districts in Asia but not Africa, and (ii) there is no association for the men in either region.

As can be seen, the slope with respect to population density is very similar for India and Asia in Figure 6a.¹⁸ The status-based motivation for the rural FLFNP-population density association, which was supported by specific consumption behaviors in India, may thus extend to the Asian region as a whole. But then why is this association absent in Africa? Our explanation for these inter-regional differences is based on the idea that the status mechanism can only be utilized to allocate resources

¹⁸While the rural FLFNP-population density association may be very similar in India and the rest of Asia, the *level* of FLFNP is much higher in India. Other non-status determinants of FLFNP, such as patriarchal norms, presumably contribute to this regional difference.

between groups in stratified (hierarchical) societies. The traditional view, both in economics (going back to Adam Smith) and anthropology (Cancian, 1976) was that social stratification in pre-modern and developing economies is positively associated with agricultural productivity. However, Mayshar et al. (2022) have recently shown that it is *not* agricultural productivity, but the type of crop that matters. In particular, the cultivation of storable cereals, as opposed to perishable roots and tubers, is a pre-requisite for the emergence of stratified societies. Although Mayshar et al. do not emphasize the Asia-Africa divide, this is evident in Figure 2 of their paper: roots and tubers are grown in abundance in Africa, whereas agriculture in Asia is restricted to the cultivation of cereals. Goody (1971) uses differences in marital arrangements to provide independent support for the inter-regional divide: his argument is that status-group endogamy, as observed in Eurasia, is a pre-requisite for a stratified social order, and this is not observed in Africa. Without social stratification, status is less useful as an allocative mechanism, which explains why females do not appear to be increasingly withdrawn from the labor force at higher levels of population density (agricultural productivity) in Africa.¹⁹

While we are able to provide a status-based explanation for inter-regional differences in the FLFP-population density association, non-status explanations are also available. For example, it is well known that the demand for female labor in agricultural production has historically been higher in Africa than in Asia due to differences in growing conditions (Boserup, 1970). These historical work patterns could have crystallized into gender norms that continue to shape female labor force participation today (Alesina et al., 2013; Bertrand et al., 2015). Population density will be increasing in agricultural productivity in both Asia and Africa. However, the lower demand for, or supply of, female labor in Asia will create a gap between the number of women in the population and in the workforce, and if this gap is increasing in agricultural productivity, then the observed cross-regional patterns could be obtained without a role for social status. The advantage of the Indian setting, which is the focus of much of our analysis, is that particular consumption behaviors can be used to buttress the claim that the status mechanism is relevant. In addition, Indian society is clearly and visibly stratified by caste. The model that follows will generate implications for variation in FLFP within districts over time, by caste group, that will be used to provide independent support for the status mechanism.

3 The Status Game

3.1 Ingredients of the Model

Our model is based on previous characterizations of social status in the economics literature; e.g. Frank (1985); Cole et al. (1992); Bagwell and Bernheim (1996); Fershtman et al. (1996); Postlewaite (1998); Weiss and Fershtman (1998). These papers, in turn, build on the seminal contributions of

¹⁹This argument is not inconsistent with recent experimental evidence that signals revealing unobserved individual income are relevant in (urban) Africa; e.g. Macchi (2023); Dupas et al. (2024). It simply says that status will be less important in a less stratified society because it cannot be used to allocate resources between groups (as in Asia).

Veblen (1899) and Weber (1922) and have the following features in common:

1. Wealth is not publicly observed and, hence, households signal their wealth by making costly visible choices; for example, by withdrawing their women from the labor force. Wealthier households have higher status in equilibrium.
2. Households have a concern for relative standing, and are willing to bear the associated signaling costs, because it is instrumental in determining their consumption of non-market goods and services; i.e. they do not necessarily value status *per se*. These instrumental concerns arise because markets are incomplete or function imperfectly.
3. The status game can be played between individuals or groups. Either way, social status is inherently relative and, hence, the allocation of non-market goods and services through this mechanism is a zero-sum game. In rural India, an individual's identity and their status is determined by their caste (Srinivas, 1967). The status game will thus be played between caste groups.

3.2 Population and Preferences

The status game is played by the local population in each village. This population consists of two (caste) groups: H and L . Each group consists of N households.²⁰ We are interested in modeling the status game between groups and, hence, all households within a group $k \in \{H, L\}$ are assumed to have the same wealth or income endowment, y_k , in a given village. The income endowments $y_k \in \{y_H, y_L\}$ vary across villages and their levels in a given village are private information; i.e. the external agents who are using the status mechanism to allocate resources do not know their value. $y_H > y_L$ in all villages and, hence, the H group always has a higher social position than the L group (the rank is fixed). However, the magnitude of this advantage will vary across villages, depending on the levels of y_H, y_L , which are revealed in equilibrium.

Households derive utility from the consumption of market goods and from a non-market good that is allocated through the status mechanism. Status is increasing with wealth, but since the wealth of household i belonging to group k , $y_{i,k}$, is unobserved, it must be signaled by a costly choice, $c_{i,k}$. This signal, in turn, determines the household's share of the non-market good, which has per capita (or, to be more precise, per household) value, v . Assuming that preferences over the consumption of market goods are logarithmic and normalizing so that the price of the consumption bundle is equal to one, household i in group $k \in \{H, L\}$ derives the following utility from consumption:

$$\log(y_{i,k} - c_{i,k}) + \frac{\mathbb{C}_k}{\mathbb{C}_k + \mathbb{C}_{-k}} \cdot 2v, \quad (1)$$

where $\mathbb{C}_k$ is the total signaling cost borne by group k and $\mathbb{C}_{-k}$ is the corresponding statistic for the other group. If incomes were observed, then $\mathbb{C}_k, \mathbb{C}_{-k}$ would be replaced by total incomes in each group, $\mathbb{Y}_k, \mathbb{Y}_{-k}$. As in Esteban and Ray's (2011) model of inter-group conflict, each household

²⁰We make this assumption for analytical convenience. Later in Section 5, when we estimate the model, we will allow group sizes to vary.

makes its signaling choice independently, but the non-market good is allocated at the level of the group, based on the aggregate contributions of all members (which determines the group's status). This connection between our model and the Esteban-Ray model is not coincidental. As discussed in the Introduction, social status and conflict are alternative (costly) mechanisms that can be used to allocated resources between groups in an economy.

Note that households could derive utility from the consumption component of the status signal. For example, if conspicuous consumption of positional goods is used as a signal, then individuals in the household might benefit from the consumption of such goods. Alternatively, if FLPNP is used as the signal, then the woman's time could be used for home production, which includes investments in children's human capital. In the Indian context, $c_{i,k}$ will also incorporate the monetized value of the non-pecuniary costs that must be simultaneously borne to achieve high status, such as vegetarianism and teetotalism. The only restriction on the signaling cost is that the household must be worse off on net; i.e. $c_{i,k}$ must be positive, for the signal to reveal its underlying wealth.

3.3 The Status Equilibrium

Household i in group k chooses its wealth signal to maximize expression (1), taking the signaling choices of the remaining households in its group and all households in the other group as given. Since all households in a group have the same income endowment, this is a symmetric equilibrium and, hence, the optimal choice of $c_{i,k}$ for $k \in \{H, L\}$ is determined by the following first-order condition:

$$\frac{1}{y_k - c_k} = \frac{c_{-k}}{(c_k + c_{-k})^2} \cdot 2 \frac{v}{N}. \quad (2)$$

This constitutes a system of two equations with two unknowns, c_H and c_L . To solve these equations, we first divide one by the other and collect terms to obtain:

$$\frac{c_k}{y_k} = \frac{c_{-k}}{y_{-k}}. \quad (3)$$

Both groups expend the same share of their income endowment on signaling in equilibrium. It follows that the non-market good would be allocated in exactly the same way if y_H , y_L were observed, although the inefficient signaling costs would not be incurred. The expenditures on signaling are strategic complements and so one implication of the status game is that an exogenous increase in c_{-k} will be accompanied by a corresponding increase in c_k . This is the approach taken by Kim et al. (2024) in their status-based analysis of educational investments in Korea. We take an alternative approach, which is more closely related to Charles et al. (2009); Bursztyn et al. (2018); Atkin et al. (2021) that exploits exogenous variation in the per capita value of status, v , and group-specific incomes, y_H and y_L .²¹

²¹Bursztyn et al. (2018) document that an exogenous increase in the value of a product that is positively associated with social status increases consumers' willingness to pay for it. Atkin et al. (2021) document that status-based investments are increasing in group income, but there is no reference group (with its associated income) in their

Proposition 1 *The average signaling cost in a local population is (i) increasing in the per capita value of status, (ii) increasing in the mean income endowment, and (iii) decreasing in the income endowment gap between the groups.*

To prove the proposition (see Appendix C for the complete derivation) we first substitute from (3) in (2) and then derive an expression for c_k as a function of the exogenous variables in the model:

$$c_k = \frac{y_k}{1 + Kw}, \quad (4)$$

where $K \equiv \frac{(y_H + y_L)^2}{y_H y_L}$ and $w \equiv \frac{N}{2v}$.

Taking the average over $k = H, L$ and denoting the average signaling cost by $\bar{c} = \frac{c_H + c_L}{2}$ and the mean income endowment by $\bar{y} = \frac{y_H + y_L}{2}$:

$$\bar{c} = \frac{\bar{y}}{1 + Kw}. \quad (5)$$

Observe that K in the denominator of equation (5) can be expressed as a function of the mean income endowment, $\bar{y}$, and the income endowment gap, $\Delta y \equiv \frac{y_H - y_L}{2}$:

$$K = \frac{4\bar{y}^2}{\bar{y}^2 - \Delta y^2}.$$

Differentiating the preceding equation, it is straightforward to verify that K is increasing in Δy since that term only appears in the denominator on the right hand side. It can also be shown that K is decreasing in $\bar{y}$ (see Appendix C). This implies, from equation (5), that $\bar{c}$ is decreasing in Δy and increasing in $\bar{y}$, since $\bar{y}$ also appears in the numerator of that equation. It also follows from equation (5) that $\bar{c}$ is increasing in the per capita value of status, since w is decreasing in v , to complete the proof.

The preceding result allows us to interpret the cross-sectional evidence presented in Section 2 through the lens of the model: v is increasing in population density (agricultural productivity) and this, in turn, leads to higher $\bar{c}$, which we measured by female labor force non-participation and associated consumption choices (vegetarianism and teetotalism). As derived below, c_L , c_H are also increasing in v , which is in line with the caste-specific results that were reported in that section.

As a corollary to Proposition 1, we can derive implications for group-specific investments in social status. From equation (4),

$$c_L = \frac{\bar{y} - \Delta y}{1 + Kw} \quad (6)$$

$$c_H = \frac{\bar{y} + \Delta y}{1 + Kw} \quad (7)$$

analysis. In Charles et al. (2009), individuals with relatively high incomes attempt to separate themselves from their own group by consuming conspicuously, but the strategic interaction between groups that is a key feature of our analysis is absent.

Differentiating equations (6) and (7), it is straightforward to verify that the qualitative implications of the model, derived in Proposition 1 for $\bar{c}$, apply to c_L , c_H as well, with one exception; the sign of the association between c_H and Δy is ambiguous. This is on account of the Δy term in the numerator of equation (7), which works against the negative effect of Δy on c_H , through the K term in the denominator, that we derived above.

If we make the additional assumption that $\frac{\Delta y}{(1+Kw)^2} \approx 0$, then the model also has testable implications for the magnitude of these effects (see Appendix C):²²

$$\frac{\partial \bar{c}}{\partial \bar{y}} = \frac{\partial c_L}{\partial \bar{y}} = \frac{\partial c_H}{\partial \bar{y}} \quad (8)$$

$$\left| \frac{\partial c_L}{\partial \Delta y} \right| > \left| \frac{\partial \bar{c}}{\partial \Delta y} \right| > \left| \frac{\partial c_H}{\partial \Delta y} \right| \quad (9)$$

We will test a linear approximation to the model in the section that follows and, hence, we do not want to take these quantitative implications too literally. With regard to (8), a more reasonable expectation is that the $\bar{y}$ effect will be of comparable magnitude with $\bar{c}$, c_L , c_H as outcomes. Intuitively, a secular increase in the income endowment for both groups, holding Δy constant, will generate a similar increase in their status signals. In addition, (9) tells us that the lower castes, who are seeking to raise their social position, will respond more to a narrowing of the income endowment gap than the high castes, who are pushing back to maintain their position. The differential increase in FLFP, as the income endowment between caste groups narrows, is a novel implication of the status model and we will use this result to buttress the claim that our findings are not being driven by income effects or changes in patriarchal norms over time.

4 Testing the Model

Estimating equation: We test the model's implications, as specified in Proposition 1, by estimating the following equation with NSS data, across districts j and over rounds or time periods t :²³

$$c_{jt} = \beta_1 \bar{y}_{jt} + \beta_2 \Delta y_{jt} + \delta_j + \gamma_t + \epsilon_{jt}. \quad (10)$$

c_{jt} denotes the signaling cost, which is measured by average or caste-specific FLFP. $\bar{y}_{jt}$ measures the mean income endowment across the two caste groups and Δy_{jt} measures the difference between the high-caste and low-caste endowments. When discussing the empirical results, we will refer to

²²This assumption will be satisfied if the quadratic term, $(1+Kw)^2$, is an order of magnitude larger than the income-gap between the two groups, Δy . However, we still need to assume that $\frac{\bar{y}}{(1+Kw)^2}$ has finite value (see Appendix C). This may not be unreasonable since $\bar{y}$ is seven times larger than Δy on average in our data.

²³The district covers a large area and has a substantial rural population, which covers many villages. The implicit assumption when we test the model at the district level is that villages are homogeneous within district-time periods. We will make the same assumption when estimating the structural model below.

the income endowments as *potential* incomes; i.e. the incomes that would be obtained if women were not withdrawn from the labor force. The district effects, δ_j , incorporate the per capita value of social status, v , as well as other fixed factors outside the model, such as patriarchal norms, that independently determine FLFP. The time-period effects, γ_t , account for secular changes that affect FLFP in all districts, while unobserved district-time period effects are captured by the ϵ_{jt} term. Note that the additive separability in equation (10) accounts for an important feature of Proposition 1, which is that the effect of each determinant of the signaling cost – $v, \bar{y}, \Delta y$ – is derived *conditional* on the other determinants. Based on that proposition, we expect $\beta_1 > 0, \beta_2 < 0$. In a rural (agrarian) economy, incomes will be derived from labor and from land (for those households who are endowed with that factor of production). We focus on wage income for the core tests of the proposition, but incorporate income from land in extensions to these tests.

To match more closely with the model, we would want to multiply FLFP by the female market wage to give us a measure of the *monetary* cost of withdrawing women from the labor market, and we will do this when estimating the structural parameters of the model in the section that follows. We omit the wage multiplier from the current analysis because its presence would undermine the validity of the instruments that we construct for $\bar{y}_{jt}$ and Δy_{jt} , as discussed below. This omission does not affect the signs of β_1 and β_2 , as implied by Proposition 1, because any factor that increases (decreases) FLFP would also increase (decrease) female wages through its general equilibrium effect.

Our model describes household decisions, whereas its implications are tested at the district level. To map the model to the data, we assume that the ‘representative’ household in each caste has two members – a male and a female – each of whom is endowed with a single unit of time (Hansen, 1985). The male devotes all his available time to work and receives the market wage, while the female’s time is allocated optimally at the intensive margin, trading off her wage income against the gain in social status when she reduces her presence in the labor market. While employment lotteries at the household level, as in Rogerson (1988) generate discrete labor market outcomes – women either enter the labor force or stay at home – the average FLFP in a given district corresponds to underlying household-level choices at the intensive margin in our model.

Variable construction: To construct the potential (labor) income terms, which appear on the right hand side of equation (10), we first measure the average wage in each district-time period at the caste (k) and gender (g) level: w_{kg} , where $k \in \{H, L\}$ and $g \in \{m, f\}$. If there was an equal share of low-caste and high-caste households in the population, as assumed in the model, and a single male and female in each household, as assumed above, then $\bar{y}$ would be constructed as an unweighted average of w_{kg} across castes and genders. In practice, castes and genders will not be balanced and, hence, we construct $\bar{y}$, and Δy , as follows in each district j and time period t :

$$\bar{y} = \sum_k x_k \sum_g x_{kg} w_{kg} \quad (11)$$

$$\Delta y = \sum_g x_{Hg} w_{Hg} - \sum_g x_{Lg} w_{Lg}. \quad (12)$$

where x_k measures the share of caste- k households, $x_H + x_L = 1$, and x_{kg} measures the share of working-age individuals by gender in each caste, $x_{km} + x_{kf} = 1$.

As in the model, the implicit assumption when constructing these statistics is that individuals are homogeneous within caste-gender subpopulations in a given district-time period. This allows us to assign the observed wage to all working-age individuals when constructing potential incomes, even if they are self employed or withdrawn from the labor force.²⁴ As discussed below, this ‘representative’ agent assumption can be relaxed once we instrument for potential incomes, but it will be retained when we estimate the model in the section that follows.

There are three potential sources of bias when $\bar{y}$, Δy are measured as above: reverse causality, omitted variables, and measurement error. We describe each source of bias below, proposing an instrumental variable strategy that addresses all of them.

Equation (10) is derived from a model in which households are making independent choices, taking the market wage (which determines their income endowment) as given. Once we aggregate up to the district level, variation in the *supply* of female labor, due to the status mechanism or unobserved factors incorporated in the ϵ_{jt} term, will affect the equilibrium female wage, which, in turn, determines $\bar{y}_{jt}$, Δy_{jt} in each district-time period. This reverse causation will arise if the production technology exhibits diminishing marginal productivity with respect to labor or if there is heterogeneity in individual ability (and selection into the labor force varies with ability). Omitted variable bias will arise if the unobserved supply-side shifters (ϵ_{jt}) are correlated with $\bar{y}_{jt}$, Δy_{jt} in equation (10) and measurement error arises because our potential income measures are based on a sample of households.

To address the potential biases listed above, we construct statistical instruments for $\bar{y}$, Δy that leverage exogenous variation in the *demand* for labor. This is the classical approach to identify the supply response to price changes (wages in our context) and will also address omitted variable bias caused by supply-side factors as well as measurement error. Our analysis is based on a sample of rural households and in an agrarian economy, the demand for labor at any point in time will depend on local contemporaneous rainfall (Jayachandran, 2006). This is true not only for individuals engaged in agriculture, but also for those employed in other occupations. We thus use rainfall, available annually at the district level from 1901 to 2018 (see Appendix A) to construct the statistical instruments, as described below.

The objective when constructing the statistical instruments is to isolate that part of the variation in $\bar{y}$, Δy that is generated by exogenous rainfall. While rainfall may affect incomes in all occupations in an agrarian economy, this effect will not be uniform. The NSS provides the “primary occupation” of each household: (i) technical, (ii) administrative, (iii) clerical, (iv) sales and services, (v) agriculture, and (vi) others. While individuals may change jobs temporarily in response to economic shocks, it is reasonable to assume that the household’s primary occupation is fixed and predetermined. The

²⁴The additional assumption is that the shadow price of labor for a self-employed individual is equal to the market wage. This implies that there is no restriction on movement between self employment and wage labor in the local economy, which appears reasonable.

first step in constructing the statistical instruments is to nonparametrically estimate the relationship between average wages, measured at the caste-gender-occupation level, and rainfall in each district-time period, after partialling out district and time period effects.²⁵ These estimates are reported in Appendix Figure C1, where we observe substantial variation by caste and occupation, separately for males and females. Predicted wages based on these estimates, $\hat{w}_{kg}$ are then used to construct instruments for $\bar{y}$, Δy in each district-time period:

$$\bar{y}_{IV} = \sum_k \bar{x}_k \sum_g \bar{x}_{kg} \hat{w}_{kg} \quad (13)$$

$$\Delta y_{IV} = \sum_g \bar{x}_{Hg} \hat{w}_{Hg} - \sum_g \bar{x}_{Lg} \hat{w}_{Lg} \quad (14)$$

where $\bar{x}_k$, $\bar{x}_{kg}$ denote district-level averages of x_k , x_{kg} computed over all time periods. This averaging accounts for the possibility that changes in the population shares x_k , x_{kg} , within a district over time, are correlated with unobserved factors that determine female labor supply, such as changes in patriarchal norms. The direct effect of $\bar{x}_k$, $\bar{x}_{kg}$ on FLFNP is, moreover, subsumed in the district fixed effects that are also included in the estimating equation. While our instruments thus address the potential sources of bias that we listed above, notice that they are correlated with female wages, w_{kf} . This is why we do not multiply FLFNP by that variable when constructing the dependent variable in equation (10).

Estimation results: Table 5, Columns 1-3 report OLS estimates of equation (10), with $\bar{c}$, c_H , c_L as the dependent variables. Table 5, Columns 4-6 report the corresponding IV estimates. As discussed above, the signalling costs are measured by FLFNP.

As implied by the model, the coefficient on mean potential income, $\bar{y}$, is positive and significant with all specifications in Table 5. The coefficient on the caste-gap in potential incomes, Δy , is negative, and significant with one exception (when the dependent variable is high-caste FLFNP). Recall that the model does not unambiguously sign this particular coefficient and, hence, this result is not unexpected.

In the corollary to Proposition 1, we derived additional implications with respect to the coefficients across caste groups: (i) We expect the magnitudes of the $\bar{y}$ coefficients to be roughly comparable. (ii) We expect to observe an ordering in the (absolute) magnitude of the Δy coefficients; the low castes should have the largest coefficient and the high castes the smallest coefficient, with the average coefficient lying in between. Focussing on the IV estimates in Columns 4-6 of Table 5, we see that the results match the more nuanced predictions of the model.

Since district fixed effects are included in both the first-stage and the second-stage equation, the instrumental variable estimates leverage variation in rainfall within districts over time to predict changes in $\bar{y}$, Δy . The variation in the demand for labor and, by extension, the market wage

²⁵Although the model is static, it is possible that rural labor markets adjust gradually to rainfall shocks and, hence, we measure rainfall in period t as a three-year average over periods $t-2$, $t-1$ and t . The results that follow (available on request) are robust to alternative construction of the rainfall variable.

Table 5: Female labor force non-participation within districts over time

Dep. variable	FLFNP					
	OLS			IV		
	all	high	low	all	high	low
Regression:	(1)	(2)	(3)	(4)	(5)	(6)
Caste group:						
mean potential income	0.254*** (0.059)	0.185*** (0.061)	0.356*** (0.075)	0.881*** (0.159)	0.873*** (0.176)	1.020*** (0.199)
caste-gap in potential income	-0.101*** (0.038)	-0.023 (0.040)	-0.217*** (0.048)	-0.274** (0.122)	-0.128 (0.126)	-0.563*** (0.180)
Kleibergen-Paap Wald F-statistic	—	—	—	85.67	85.67	85.67
Kleibergen-Paap LM statistic	—	—	—	82.68	82.68	82.68
Dep. var. mean	0.648	0.683	0.582	0.648	0.683	0.582
Observations	2832	2832	2832	2832	2832	2832

Source: NSS (“thick” and “thin” rounds)

District and NSS round fixed effects are included in the estimating equation.

Standard errors are clustered at the level of 1981 district boundaries.

that is generated by the rainfall shocks is independent of labor supply shifters, including the status mechanism and patriarchal norms (which would appear in the ϵ_{jt} term in equation (10)). While our instruments thus plausibly satisfy the exclusion restriction, they must also have sufficient statistical power to be valid. Based on the first-stage regressions reported in Appendix Table C1, we do not face a weak instrument problem and the Kleibergen-Paap F statistic in Table 5 is above 80. Since we have one source of exogenous variation – rainfall – and two endogenous variables – $\bar{y}$, Δy – one remaining requirement for our two instruments to be valid is that rainfall should have a differential effect on wages by caste. Reassuringly, the Keibergen-Paap LM statistic, which tests for under-identification is also above 80 in Table 5. We complete the tests of the model by verifying the robustness of the results in various ways in Appendix C.

1. We include eight “thick” and “thin” NSS rounds in the sample in Table 5 to increase variation over time. As a robustness check, we only include the five “thick” rounds in Appendix Table C2.

2. We replace the district-level averages, $\bar{x}_k$, $\bar{x}_{kg}$ by the corresponding time-period averages when constructing the instruments in Appendix Table C3. The time-period effects that are also included in the estimating equation will now subsume the direct effect of the national-level population shares on FLFNP. While there are hundreds of districts, there are only eight time periods (NSS rounds) in our sample. There is consequently much less variation in these instruments relative to the benchmark specification in Table 5. Nevertheless, the $\bar{y}$, Δy coefficients retain their statistical significance and are very similar in magnitude to the point estimates in that table.

3. The per capita value of status, v , is subsumed in the district fixed effect in equation (10). However, it is possible that v changes over time with economic development. As the economy grows

larger, the value of status will increase, following the same argument that we used to motivate a positive association between v and population density (agricultural productivity) in the cross-section. At the same time, markets will expand with economic development, with an accompanying decline in the need for the social status mechanism. While the nature of the variation in v over time within a district is thus theoretically ambiguous, we allow for such variation by including population density interacted with round (time period) effects in the estimating equation. The results with this augmented specification of equation (10) are reported in Appendix Table C4, where we see, once again, that the point estimates are very similar to the corresponding estimates with the benchmark specification in Table 5.

4. While the tests of the model thus far have focused on wages, income from land will also be relevant in a rural economy. If land ownership and productivity were available by caste in each district-time period, then we could construct measures of $\bar{y}$, Δy based on land incomes and test the model independently. However, district-level information on land in the NSS is restricted to the Land and Livestock Holding Survey, conducted in the 2003 round, which lists the amount of land owned by each caste group. Without information on land productivity, and time varying data more generally, we cannot independently test the model. Nevertheless, we would like to control for land incomes in the estimating equation because they will vary with rainfall (our instrument). Land markets are extremely thin in India and it is thus reasonable to assume that land holdings are fixed over time. For each caste group k , the ‘representative’ household’s income from land in district j and time period t can then be parsimoniously specified as $\gamma_k R_{jt} \frac{A_{jk}}{N_{jk}}$, where the γ_k parameters (to be estimated) measure caste-specific land productivity, R_{jt} is rainfall, A_{jk} is total acreage owned by caste k in district j in 2003, and N_{jk} is the number of households in that caste in that district in that year (obtained from the NSS Employment and Unemployment Survey). The robust results with an augmented specification of equation (10) that includes the land income terms are reported in Appendix Table C5.

In this section and in Section 2 we have provided empirical support for the claim that social status considerations determine FLFP in India. Our tests are based on independent sources of variation: (i) cross-sectional variation in population density (agricultural productivity) across districts, and (ii) variation in *potential* incomes within districts over time, generated by rainfall shocks. The positive coefficient on mean potential income is consistent with the status model, but could also be generated by an income effect. However, it is difficult to come up with a non-status explanation for why FLFP *increases* as the gap between low-caste and high-caste potential incomes *narrows*. It is even more difficult to explain why the magnitude of the coefficients on mean potential income and the caste-gap in potential income, by caste, match the more nuanced implications of the model.

5 Quantitative Analysis

The empirical tests thus far have been based on the qualitative implications of the model. The next step in the analysis is to estimate its structural parameters. We will then use the estimated model to (i) assess its ability to *predict* variation in FLFNP in each NSS round, (ii) explain why FLFNP increased over time, despite the increase in female education, as documented in Figure 1, and (iii) examine alternative policies that could potentially reduce FLFNP.

In our analytical model, households choose the signaling cost, which we measure in practice by FLFNP, taking potential incomes (and market wages) as given. When testing the implications of this model, we accounted for the reverse effect of FLFNP on wages by instrumenting for them. For the counter-factual policy analysis mentioned above, we will want to allow for general equilibrium effects and so wages will be endogenized in the structural model that follows. Since we are also interested in making sense of the positive association between FLFNP and female education, we will add education choice to the structural model.

5.1 Structural Estimation

Set up of the model: As in the analytical model, the household consists of a male and a female member, each of whom is endowed with a single unit of time. In the augmented structural model, household i belonging to caste k can allocate each member's time to skilled tasks, $\xi_{i,k}$, or unskilled tasks, $(1 - \xi_{i,k})$. Skilled tasks require investments in education, which cost $e_{kg}(\xi_{i,k})$.²⁶ The household's potential income can then be expressed as follows:

$$y_{i,k} = \sum_g w_{sg} \xi_{i,k} + w_{ug}(1 - \xi_{i,k}) - e_{kg}(\xi_{i,k}) \quad (15)$$

where w_{sg} , w_{ug} are the wages faced by the household, which vary by skill, caste, and gender in each district-time period.

The household's expenditure on status signaling, by skill, can be expressed as follows:

$$c_{i,sk} = w_{skf} \xi_{i,k} \tau_{i,sk} \eta_{sk} \quad (16)$$

$$c_{i,uk} = w_{ukf}(1 - \xi_{i,k}) \tau_{i,uk} \eta_{uk}, \quad (17)$$

where $\tau_{i,sk}$, $\tau_{i,uk}$ are the fractions of the skilled-task time and the unskilled-task time that are withdrawn from the labor market for the female member of household i . The η_{sk} , η_{uk} parameters, both of which are less than one, reflect the idea that some of the lost wage income is recouped by the household because a woman who is withdrawn from the labor market could contribute to home production. These parameters also implicitly account for patriarchal norms and other cultural fac-

²⁶We allow the cost of education to vary by ethnicity (caste in this case) and gender, as in Hsieh et al. (2019). This generates differences in education levels, by caste and gender, as observed in our data.

tors that restrict female labor force participation in developing countries (Jayachandran, 2015) all of which effectively dampen the cost to the household of withdrawing the woman from the labor force.²⁷ Notice that η_{sk} , η_{uk} vary by skill and caste group, since an educated woman could contribute more to home production (child rearing) and patriarchal norms vary by caste status, as discussed in Section 2.

Given its potential income, as specified in equation (15), household i chooses $\xi_{i,kg}$, $\tau_{i,sk}$, $\tau_{i,uk}$ to maximize the following objective function:

$$\log(y_{i,k} - c_{i,sk} - c_{i,uk}) + \frac{\mathbb{C}_k}{\mathbb{C}_k + \mathbb{C}_{-k}} \cdot 2v. \quad (18)$$

This function is analogous to expression (1), except that the household separately chooses how much time to withdraw from the female's skilled-task time endowment and unskilled-task time endowment. This determines $c_{i,sk}$ and $c_{i,uk}$ from equations (16) and (17), which, in turn, enter the status function in equation (18).²⁸ Although education decisions are made before labor market decisions, there is no uncertainty in the model and, hence, the optimal $\xi_{i,kg}$, $\tau_{i,sk}$, $\tau_{i,uk}$ can be determined simultaneously from the first-order conditions that are derived from this maximization problem. Since all households in a given caste group are identical in each district-time period, we can derive expressions for ξ_{kg} , τ_{sk} , τ_{uk} as (implicit) functions of caste-gender specific wages and the parameters of the model from these first-order conditions, just as we did when solving the analytical model (see Appendix D). As with the analytical model, these choices at the intensive margin map into district-time period level outcomes once we introduce *ex post* lotteries: ξ_{kg} maps into the fraction of educated individuals by caste-gender and $\xi_{kf}\tau_{sk} + (1 - \xi_{kf})\tau_{uk}$ maps into FLFNP by caste. However, these outcomes will also have a reverse effect on wages, which thus cannot be treated as exogenous, and hence the next step is to derive expressions for wages at the skill-caste-gender level.

For the purpose of the structural model, we assume that the status game is played in each village between a finite number of households, while the wage is determined competitively at the district level. Each district consists of a large number of homogeneous villages. While there are an equal number of low-caste (L) and high-caste (H) households, N , in the analytical model, we now allow these numbers to vary: each village in a given district and time period has a fraction $x_L = \frac{N_L}{N_L + N_H}$ low-caste households and there is a corresponding mass x_L of low caste households in that district, with $x_H \equiv 1 - x_L$. We assume that output in the district-time period is determined by a linear aggregate production function: $Y = AE$, where A is total factor productivity and E is aggregate labor. Labor is heterogeneous along three dimensions: gender, caste, and skill. We thus use a nested-CES structure, as in Card and Lemieux (2001), Ottaviano and Peri (2012), to aggregate the different

²⁷ η_{sk} , η_{uk} will also incorporate the monetary-equivalent costs associated with behaviors such as vegetarianism and teetotalism that are needed, together with FLFNP, to achieve high status in India. This will amplify the monetary cost of removing the woman from the labor force and, hence, the implicit assumption is that this channel is dominated by the factors listed above that dampen these costs.

²⁸As described in Appendix D, we allow $c_{i,sk}$ and $c_{i,uk}$ to be imperfect substitutes or even complements in the status function (although they could enter additively as a special case).

components of labor:

$$\begin{aligned}
E &= \left[\theta_f E_f^\rho + \theta_m E_m^\rho \right]^{\frac{1}{\rho}} \\
E_g &= \left[\theta_{Lg} E_{Lg}^{\rho_g} + \theta_{Hg} E_{Hg}^{\rho_g} \right]^{\frac{1}{\rho_g}}, \quad g = \{f, m\} \\
E_{kg} &= \left[\theta_{skg} E_{skg}^{\rho_{kg}} + \theta_{ukg} E_{ukg}^{\rho_{kg}} \right]^{\frac{1}{\rho_{kg}}}, \quad k = \{H, L\} \\
E_{skf} &= \xi_{kf}(1 - \tau_{sk})x_k, \quad E_{ukf} = (1 - \xi_{kf})(1 - \tau_{uk})x_k \\
E_{skm} &= \xi_{km}x_k, \quad E_{ukm} = (1 - \xi_{km})x_k
\end{aligned}$$

The labor productivity parameters vary by caste and gender, conditional on skill, in the preceding specification. This could be due to discrimination by ethnicity (caste) and gender, as also assumed by Hsieh et al. (2019). Alternatively, this could reflect an identity-based preference for caste-specific traditional occupations (Cassan et al., 2021; Oh, 2023) or the presence of caste networks in particular, not necessarily traditional, occupations (Munshi, 2019). While we thus allow for market frictions, the assumption is that labor is allocated efficiently within a district-time period, conditional on the differences in productivity.²⁹ The wage for each skill-caste-gender category is thus determined by the associated marginal productivity of labor:

$$w_{skg} = \frac{\partial Y}{\partial E_{skg}} = \frac{\partial Y}{\partial E} \times \frac{\partial E}{\partial E_g} \times \frac{\partial E_g}{\partial E_{kg}} \times \frac{\partial E_{kg}}{\partial E_{skg}}, \quad (19)$$

$$w_{ukg} = \frac{\partial Y}{\partial E_{ukg}} = \frac{\partial Y}{\partial E} \times \frac{\partial E}{\partial E_g} \times \frac{\partial E_g}{\partial E_{kg}} \times \frac{\partial E_{kg}}{\partial E_{ukg}}. \quad (20)$$

As shown in Appendix D, skilled and unskilled wages, by caste and gender, can be derived as functions of the education investment, ξ_{skg} , and the time withdrawn from the skilled-task and unskilled-task endowments, τ_{sk} , τ_{uk} , from equations (19) and (20) to close the model.

Solving and estimating the model: The first-order conditions with respect to ξ_{kg} , τ_{sk} , τ_{uk} provide us with four equations for each caste group k . As noted, if wages, w_{skg} , w_{ukg} , were treated as exogenous, then we could solve these equations simultaneously to compute the equilibrium ξ_{kg} , τ_{sk} , τ_{uk} in each district-time period. Since we endogenize wages as well, however, the model must be solved iteratively, using the following algorithm:

Step 1. Guess w_{skg} , w_{ukg}

Step 2. Given w_{skg} , w_{ukg} from Step 1, solve for ξ_{kg} , τ_{sk} , τ_{uk}

Step 3. Given ξ_{kg} , τ_{sk} , τ_{uk} derived in Step 2, solve for w_{skg} , w_{ukg} from equations (19) and (20).

²⁹Hsieh et al. (2019) introduce a wedge between marginal productivity and the realized wage that is ethnicity (race) and gender specific, while assuming that productivity is the same in all groups. Regardless of the way in which distortions are introduced, market clearing wages will vary by ethnicity and gender in equilibrium.

Use w_{skg} , w_{ukg} derived from Step 3 as the guess in Step 1 for the next iteration and continue to iterate in this way until there is convergence; i.e. the guess in Step 1 matches the wages derived in Step 3.

The algorithm described above allows us to solve the 16 endogenous variables in the model, given its parameters. To estimate these parameters, we focus on the five “thick” NSS rounds: 1987, 1999, 2003, 2009, and 2011. In each round, we construct a (predicted) log population density grid, such that each grid interval contains an equal number of districts. The number of intervals is set to 10. In each interval, we compute (i) mean FLFP, by skill and caste, (ii) mean education, by skill, caste, and gender, and (iii) mean wages, by skill, caste, and gender. This leaves us with 160 data moments in each survey round.³⁰ For a given set of structural parameters, we can solve the 16 endogenous variables in the model in each interval and survey round, as described above, which then allows us to compute the model moments that correspond to the data moments. We search over all parameter values, using an algorithm described in Appendix D, to find the set of parameters that minimizes the (percentage) difference between the data moments and the model moments. There are 37 structural parameters, listed in Appendix Table D1, and 160 moments for matching, leaving us with sufficient degrees of freedom for estimation in each survey round.


5.2 Model Fit

We match 160 data moments and model moments, as closely as possible, when estimating the structural parameters in each survey round. With five survey rounds, this leaves us with 800 moments. Given the large number of moments, we first follow Oswald (2019), Heise and Porzio (2022), and take a graphical approach to report the model fit. Figure 7 plots that model moment on the y axis that corresponds to each data moment (on the x axis). If the moments match perfectly, then all points would lie on the 45 degree line. Figure 7a reports the goodness of fit for the FLFP moments, separately by caste. Figure 7b reports the corresponding graph for the education and wage moments, combining castes and genders. For completeness, we report the education and wage moments separately by caste and gender in Appendix Figures D1 and D2. We see in all figures that the points are tightly clustered around the 45 degree line. Despite the model’s parsimonious structure – we estimate 185 parameters by targeting 800 moments across all survey rounds – it still fits the data very well.

While the graphical approach allows us to include all targeted moments in the figures that we present, it does not tell us how the model fit varies in the cross-section with respect to population density or over time (across survey rounds). Figure 8a reports the association between FLFP and population density in the first (1987) and the last (2011) NSS rounds. This figure corresponds to Figure 3a, except that we measure mean FLFP and population density in wider intervals consisting

³⁰The number of intervals we have chosen trades off two considerations: as the number of intervals increases, we will pick up finer grained variation in the data, but the precision of our estimated data moments will also decline. While we put more weight on the second consideration by selecting a relatively small number of intervals (10), we note that the results that follow are robust to using a larger number of intervals (20). These results are available from the authors on request.

Figure 7: Comparing model and data moments




Source: NSS, 1951 population census

of multiple districts, rather than at the district level. We see that the qualitative patterns in Figure 3a are retained in Figure 8a. Moreover, the FLFNP predicted by the model matches closely with the data moments, across the range of population densities in each survey round. Figure 8b plots the change in FLFNP over time (averaged across all districts) and we see that the model also matches the time trend very closely. For completeness, Appendix Figures D3 - D6 examine the fit of the model with respect to male and female education and wages, just as we did with FLFNP in Figure 8. We see that education and wages increase steeply over time and that our model can fit these trends extremely well.

We have ignored land income and assumed that household incomes are determined by wages alone when estimating the model thus far. Nevertheless, and despite its parsimonious structure, the model does an excellent job of fitting the data. We thus would not expect much improvement on this dimension if income from land was incorporated in the structural model. The concern that remains, however, is that inclusion of the land income component could change other estimated parameters of the model and this would then affect the results of the counterfactual analysis in the section that follows.

We address the preceding concern by including land as a factor of production, in addition to labor, in an augmented aggregate production function, which is now characterized by a Cobb-Douglas technology. The estimation proceeds as above, with the same set of equations, except that the representative household's potential income now includes a land component (see Appendix D). As

Figure 8: FLFNP: comparing the model and the data



Source: NSS, 1951 population census


noted, the NSS only provides district-level information on land in the 2003 round and, hence, the augmented model is estimated in that round alone, in two ways: (i) Caste-specific land holdings are included in the aggregate production function, using a nested-CES structure as above. (ii) Since land holdings are available separately for irrigated and unirrigated land, we also estimate a more flexible nested-CES specification with caste (high, low) and land-type (irrigated, unirrigated) as components. We see in Appendix Figure D7a that the magnitude of some parameters does change when land income is incorporated in the model. However, the counterfactual simulations described below in Figure 10 are qualitatively unchanged. As reported in Appendix Figure D7b, a decline in the cost of female education increases FLFNP, whereas an increase in the η parameters works in the opposite direction.

5.3 Parameter Estimates and Counterfactual Simulations

Our model is able to predict variation in FLFNP, education, and wages in the cross-section and over time. We complete the analysis by (i) assessing whether the parameter estimates match a key assumption of the status model and the cross-sectional analysis in Section 2, (ii) by isolating the channels through which status changes FLNP over time, and (iii) by examining alternative policies that could potentially be used to reduce FLFNP.

When we estimate the structural model, we specify that total factor productivity, A , and the value of status, v , are flexible power functions of population density (each with two parameters).

Figure 9: Value of status and status-based determinants of FLFNP



Once the parameters of the model are estimated, separately in each NSS round, v can be computed in each population density interval and time period. Recall from Section 2 that population density is measured in 1951 and that the assumption in the cross-sectional analysis is that this variable is positively correlated with v in all NSS rounds. Although we cannot measure the value of status directly, our model-based estimates of v allow us to verify this assumption. As can be seen in Figure 9, our measure of v is indeed increasing in 1951 population density in each round. Note that this result does not follow mechanically from the observation that FLFNP is increasing in population density and the structure of the model (which implies that FLFNP is increasing in v). There are many other factors in the model that could generate a positive association between FLFNP and population density, and we have also allowed the v function to have a positive or negative association with respect to population density.

Based on our model, the increase in FLFNP over time that is observed in Figure 8b can be explained by three independent status-based determinants: (i) an increase in the value of status, v , (ii) an increase in mean potential income, $\bar{y}$, and (iii) a decline in the caste-gap in potential income, Δy . We plot the change in each of these determinants over time in Figure 9b. As can be seen, the increase in FLFNP is driven by a corresponding increase in v , which is also evident in Figure 9, and by an increase in $\bar{y}$. While the remaining determinant, Δy , was especially useful for testing the model in Section 4, it does not contribute to the aggregate change in FLFNP over time.

The status-based factors underlying the increase in FLFNP that we have uncovered are natural consequences of economic development: incomes ($\bar{y}$) will increase, and so will the competition for

increasingly valuable amenities (v) as an economy grows. In the long run, markets will thicken and expand, and the status mechanism will ultimately be less relevant. In the interim period, however, it is important to implement policies that will reduce FLFP in an environment where an underlying status motivation is present and we next use the estimated model to examine such policies.


The first policy that we consider encourages women to enter the labor force by reducing the cost of education for them. Figure 10a reports the average female cost of education, combining both caste groups, in each survey round. We see that this cost has been declining over time. Our counterfactual simulations, reported in Figure 10b, reduce the estimated cost in the last (2011) NSS round by an additional 20%. Instead of reducing FLFP, we see that there is a substantial *increase* at (almost) all population density levels. Viewed through the lens of the model, this seemingly anomalous increase can be explained by the fact that the decline in the cost of education, together with the accompanying increase in education with its higher wages, would have increased *potential* incomes. This, in turn, would have increased investments in the status game and, hence, signaling costs, which we measure by FLFP. Circling back to Figure 1, which motivated our analysis, the decline in the cost of education that we have just documented would have increased female education, as observed. The resulting increase in potential income over time could have increased FLFP even further.³¹

The preceding discussion tells us that conventional policy prescriptions to increase female labor force participation, such as investments in female education, might backfire in developing economies where status considerations are relevant. This is also true for related incentive schemes that, for example, would give a monetary bonus to women, over and above the market wage, if they entered the work force. The resulting increase in their households' potential incomes, with the accompanying increase in status investments, could potentially more than offset the direct positive effect on female labor force participation. Consistent with this argument, recent experimental evidence from (urban) India indicates that female labor supply is surprisingly unresponsive to wages (Jalota and Ho, 2024). Based on our model, the way to get FLFP to decline is to make status investments more costly, per unit of female time withdrawn from the labor force, while leaving potential incomes unchanged. This can be accomplished by *increasing* the η parameters, which incorporate the returns to home production (child rearing) as well as patriarchal norms that constrain female labor force participation.

Figure 10a reports the average value of the η parameters, combining both caste groups, in each NSS round. These parameters have been increasing over time, possibly because the traditional patriarchal norms have been weakening with economic development, but they are still far below one. Figure 10b reports the counterfactual level of FLFP in the last (2011) NSS round that is generated by an additional 20% increase in the estimated η parameters in that year and we see that there is a substantial *decline* in FLFP at all levels of population density. In recent years, a number of research studies have examined interventions that are designed to target patriarchal norms and

³¹Hnatkovska et al. (2012) document, using NSS data, that education levels for low castes and high castes have converged over time. Our estimates of the cost of education also reveal such convergence (not reported). As implied by our model and verified in Section 4, the resulting convergence in potential incomes between the caste groups would also have (independently) increased FLFP over time.

Figure 10: Counterfactual policy simulations



Source: NSS, 1951 population census

other constraints to female labor force participation. Viewed through the lens of our model, these interventions make withdrawal from the labor force more costly, from the household's perspective, if they are successful. This will reduce FLFNP in equilibrium, with the status channel amplifying this effect, and this is precisely what we observe in the counterfactual simulation.³²

6 Conclusion

The widespread withdrawal of women from the labor force is a major source of inefficiency in developing economies. Take the Indian economy, for example, which is the focus of our analysis. This economy has been growing since the middle of the twentieth century, and particularly rapidly since the early 1990's. During this time, rural female labor force non-participation (FLFNP), which was high to begin with, has continued to increase, with no sign of a tapering off. Our status-based explanation for this observation is based on the idea that caste groups can signal their wealth, and thereby increase their status, by withdrawing their women from the labor force. If the value of social status or the willingness to bear the signaling cost is increasing with economic development, then this would explain the persistent increase in FLFNP.

³²Soft-touch interventions that provide information do not appear to have a substantial or sustained effect on female labor force participation (Dean and Jayachandran, 2019; McKelway, 2023). However, recent experimental evidence indicates that a two-step process in which jobs are first offered in-home, allowing gender norms to weaken, after which work outside the home is made available, may be more effective (Ho et al., 2023).

While we have a plausible status-based explanation for the dynamics of rural FLFNP, this explanation is only valid if there is a link between FLFNP and status. The first source of cross-sectional variation that we use to establish this link builds on the argument that the value of status will be greater in a large local economy. The size of a local economy is increasing in agricultural productivity, which, in turn, is proxied by population density at early stages of economic development. As expected, rural FLFNP is increasing in population density, instrumented by exogenous crop suitability, across Indian districts. The same positive association is observed more generally in Asia, but is absent in sub-Saharan Africa.

While we are able to provide a status-based rationale for the inter-regional differences that we have uncovered, which relies on the absence of vertical stratification in African society, alternative explanations are available. The advantage of focusing on India’s caste-based society in our analysis is that specific behaviors, such as vegetarianism and teetotalism, are also needed to achieve high status. Matching the FLFNP results across Indian districts, these behaviors are seen to be increasing in population density, both for the low castes and for the high castes. While cross-sectional variation in FLFNP, vegetarianism, and teetotalism for the high castes could reflect underlying variation in their (caste) identity, the only possible explanation for the low castes is status-based emulation, since these behaviors are not traditionally associated with their own identity.

The second source of variation that we use to link rural FLFNP to social status exploits exogenous variation in low-caste and high-caste potential incomes, within districts over time. Using rainfall shocks to construct statistical instruments, an exogenous increase in mean potential income and an exogenous narrowing of the caste-gap in potential incomes both generate an increase in FLFNP, as implied by our status model. This last result is especially useful in ruling out alternative explanations and also distinguishes our model from previous analyses of status, which do not examine the interaction between groups. Our model and its implications are actually more closely related to models of group conflict that have been proposed in the literature. Status and conflict are alternative ways of allocating resources when markets are missing and when ethnic groups are not vertically stratified, as in sub-Saharan Africa, then conflict will be used for this purpose.

Having documented a link between rural Indian FLFNP and status, in the cross-section across districts and within districts over time, we complete the analysis by (i) estimating the structural parameters of the model, (ii) by using the estimated parameters to disentangle the status-based factors that led to the increase in FLFNP over time, and (iii) by conducting counterfactual policy simulations.

Despite its parsimonious structure, our model fits the data very well. Based on the estimated parameters, the observed increase in FLFNP over time is largely driven by underlying increases in the value of status and (mean) potential income. Changes in these factors are a natural consequence of economic development and, hence, it is important to design policies that will reduce FLFNP in an environment where status considerations are likely to remain relevant for the foreseeable future. The first policy simulation that we consider is based on an exogenous reduction in the cost of education and we find that this *increases* FLFNP. Viewed through the lens of our model, this is because

potential household incomes increase and this, in turn, increases the competition for social status with its associated signaling costs. The more general message is that any incentive-based policy, such as a monetary bonus for women who work, that raises potential incomes could backfire in an economy where status considerations are relevant. The rapid increase in female education over time, which is a noteworthy feature of Indian economic development, could paradoxically have increased FLFP even further. The second simulation that we consider is based on a policy that reduces the non-pecuniary constraints to female labor force participation; for example, by weakening patriarchal norms. This effectively increases the cost of withdrawing women from the workforce, without changing potential incomes, and results in a substantial *decline* in FLFP. This promising result provides an additional status-based rationale for attempts to weaken patriarchal norms in developing economies.

References

- AFRIDI, F., T. DINKELMAN, AND K. MAHAJAN (2018): “Why are fewer married women joining the work force in rural India? A decomposition analysis over two decades,” *Journal of Population Economics*, 31, 783–818.
- AGTE, P. AND A. BERNHARDT (2023): “The Economics of Caste Norms: Purity, Status, and Women’s Work in India,” *Typescript*.
- AHMAD, I. (1967): “The Ashraf and Ajlaf categories in Indo-Muslim society,” *Economic and Political Weekly*, 887–891.
- ALESINA, A., P. GIULIANO, AND N. NUNN (2013): “On the origins of gender roles: Women and the plough,” *The Quarterly Journal of Economics*, 128, 469–530.
- ARDIA, D., K. BOUDT, P. CARL, K. MULLEN, AND B. G. PETERSON (2011): “Differential evolution with DEoptim: an application to non-convex portfolio optimization,” *The R Journal*, 3, 27–34.
- ASHER, S., T. LUNT, R. MATSUURA, AND P. NOVOSAD (2021): “Development research at high geographic resolution: An analysis of night-lights, firms, and poverty in India using the SHRUG open data platform,” *The World Bank Economic Review*, 35.
- ASHRAF, Q. AND O. GALOR (2011): “Dynamics and stagnation in the Malthusian epoch,” *American Economic Review*, 101, 2003–41.
- ATKIN, D., E. COLSON-SIHRA, AND M. SHAYO (2021): “How do we choose our identity? A revealed preference approach using food consumption,” *Journal of Political Economy*, 129, 1193–1251.
- BAGWELL, L. S. AND B. D. BERNHEIM (1996): “Veblen effects in a theory of conspicuous consumption,” *The American Economic Review*, 349–373.
- BASU, A. M. (1992): *Culture, the status of women, and demographic behaviour: Illustrated with the case of India.*, Clarendon press.
- BERTRAND, M., E. KAMENICA, AND J. PAN (2015): “Gender identity and relative income within households,” *The Quarterly Journal of Economics*, 130, 571–614.

- BOSERUP, E. (1970): *Woman's role in economic development*, George Allen and Unwin Ltd.
- BURSZTYN, L., B. FERMAN, S. FIORIN, M. KANZ, AND G. RAO (2018): "Status goods: Experimental evidence from platinum credit cards," *The Quarterly Journal of Economics*, 133, 1561–1595.
- CANCIAN, F. (1976): "Social stratification," *Annual Review of Anthropology*, 5, 227–248.
- CARD, D. AND T. LEMIEUX (2001): "Can falling supply explain the rising return to college for younger men? A cohort-based analysis," *The Quarterly Journal of Economics*, 116, 705–746.
- CARRANZA, E. (2014): "Soil endowments, female labor force participation, and the demographic deficit of women in India," *American Economic Journal: Applied Economics*, 6, 197–225.
- CASSAN, G., D. KENISTON, AND T. KLEINEBERG (2021): "A division of laborers: Identity and efficiency in India," Tech. rep., National Bureau of Economic Research.
- CHAKRAVARTI, U. (1993): "Conceptualising Brahmanical patriarchy in early India: Gender, caste, class and state," *Economic and Political Weekly*, 579–585.
- CHARLES, K. K., E. HURST, AND N. ROUSSANOV (2009): "Conspicuous consumption and race," *The Quarterly Journal of Economics*, 124, 425–467.
- COLE, H. L., G. J. MAILATH, AND A. POSTLEWAITE (1992): "Social norms, savings behavior, and growth," *Journal of Political economy*, 100, 1092–1125.
- DEAN, J. T. AND S. JAYACHANDRAN (2019): "Changing family attitudes to promote female employment," in *AEA Papers and Proceedings*, American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203, vol. 109, 138–142.
- DUPAS, P., M. FAFCHAMPS, AND L. HERNANDEZ-NUNEZ (2024): "Keeping up appearances: An experimental investigation of relative rank signaling," Tech. rep., National Bureau of Economic Research, Inc.
- ESTEBAN, J. AND D. RAY (2011): "Linking conflict to inequality and polarization," *American Economic Review*, 101, 1345–1374.
- FERSHTMAN, C., K. M. MURPHY, AND Y. WEISS (1996): "Social status, education, and growth," *Journal of Political Economy*, 104, 108–132.
- FRANK, R. H. (1985): "The demand for unobservable and other nonpositional goods," *The American Economic Review*, 75, 101–116.
- GALOR, O. AND Ö. ÖZAK (2016): "The agricultural origins of time preference," *American Economic Review*, 106, 3064–3103.
- GENICOT, G. AND D. RAY (2017): "Aspirations and inequality," *Econometrica*, 85, 489–519.
- GOLDIN, C. (1994): "The U-shaped female labor force function in economic development and economic history," *National Bureau of Economic Research*.
- GOODY, J. (1971): "Class and marriage in Africa and Eurasia," *American Journal of Sociology*, 76, 585–603.

- HANSEN, G. D. (1985): “Indivisible labor and the business cycle,” *Journal of Monetary Economics*, 16, 309–327.
- HARRIS, I., T. J. OSBORN, P. JONES, AND D. LISTER (2020): “Version 4 of the CRU TS monthly high-resolution gridded multivariate climate dataset,” *Scientific data*, 7, 109.
- HEATH, R. AND S. JAYACHANDRAN (2017): “The causes and consequences of increased female education and labor force participation in developing countries,” in *The Oxford Handbook of Women and the Economy*, Oxford University Press, 345–367.
- HEISE, S. AND T. PORZIO (2022): “Labor misallocation across firms and regions,” Tech. rep., National Bureau of Economic Research.
- HNATKOVSKA, V., A. LAHIRI, AND S. PAUL (2012): “Castes and labor mobility,” *American Economic Journal: Applied Economics*, 4, 274–307.
- HO, L., S. JALOTA, AND A. KARANDIKAR (2023): “Bringing work home: Flexible work arrangements as gateway jobs for women in West Bengal,” Tech. rep., Typescript.
- HSIEH, C.-T., E. HURST, C. I. JONES, AND P. J. KLENOW (2019): “The allocation of talent and U.S. economic growth,” *Econometrica*, 87, 1439–1474.
- JALOTA, S. AND L. HO (2024): “What works for her? How work-from-home jobs affect female labor force participation in urban India,” *Typescript*.
- JAYACHANDRAN, S. (2006): “Selling labor low: Wage responses to productivity shocks in developing countries,” *Journal of Political Economy*, 114, 538–575.
- (2015): “The roots of gender inequality in developing countries,” *Annual Review of Economics*, 7, 63–88.
- JUDGE, P. S. (2002): “Religion, caste, and communalism in Punjab,” *Sociological Bulletin*, 51, 175–194.
- KIM, S., M. TERTILT, AND M. YUM (2024): “Status externalities in education and low birth rates in Korea,” *American Economic Review*, 114, 1576–1611.
- KLASEN, S. (2019): “What explains uneven female labor force participation levels and trends in developing countries?” *The World Bank Research Observer*, 34, 161–197.
- LUKE, N. AND K. MUNSHI (2011): “Women as agents of change: Female income and mobility in India,” *Journal of Development Economics*, 94, 1–17.
- MACCHI, E. (2023): “Worth your weight: Experimental evidence on the benefits of obesity in low-income countries,” *American Economic Review*, 113, 2287–2322.
- MAYSHAR, J., O. MOAV, AND L. PASCALI (2022): “The origin of the state: Land productivity or appropriability?” *Journal of Political Economy*, 130, 1091–1144.
- McKELWAY, M. (2023): “Information, norms, and female employment: An experiment in India,” Tech. rep., Typescript.

- MITRA, A. AND D. RAY (2014): “Implications of an economic theory of conflict: Hindu-Muslim violence in India,” *Journal of Political Economy*, 122, 719–765.
- MUNSHI, K. (2019): “Caste and the Indian economy,” *Journal of Economic Literature*, 57, 781–834.
- OH, S. (2023): “Does identity affect labor supply?” *American Economic Review*, 113, 2055–2083.
- OLIVETTI, C. (2013): “The female labor force and long-run development: the American experience in comparative perspective,” Tech. rep., National Bureau of Economic Research.
- OSWALD, F. (2019): “The effect of homeownership on the option value of regional migration,” *Quantitative Economics*, 10, 1453–1493.
- OTTAVIANO, G. I. AND G. PERI (2012): “Rethinking the effect of immigration on wages,” *Journal of the European Economic Association*, 10, 152–197.
- POSTLEWAITE, A. (1998): “The social basis of interdependent preferences,” *European Economic Review*, 42, 779–800.
- ROBINSON, P. M. (1988): “Root-N-consistent semiparametric regression,” *Econometrica*, 931–954.
- ROGERSON, R. (1988): “Indivisible labor, lotteries and equilibrium,” *Journal of Monetary Economics*, 21, 3–16.
- SRINIVAS, M. N. (1956): “A note on Sanskritization and Westernization,” *The Far Eastern Quarterly*, 15, 481–496.
- (1967): *Social Change in Modern India*, University of California Press.
- (1977): “The changing position of Indian women,” *Man*, 221–238.
- VEBLEN, T. (1899): *The Theory of the Leisure Class: An Economic Study of Institutions*.
- WEBER, M. (1922): *Economy and Society*, University of California Press Berkeley [translated and reprinted 1978].
- WEISS, Y. AND C. FERSHTMAN (1998): “Social status and economic performance: A survey,” *European Economic Review*, 42, 801–820.

Online Appendix

A Variable Construction

1. Population Density: For the analyses with Indian (NSS, IHDS, DHS) data, we use population densities obtained from the 1951 population census, but keep track of the partitioning of districts that occurred over time. For example, if district A was divided into two districts, B and C at time t , then we measure all outcomes at the level of the contemporaneous district; i.e. based on the original district A boundaries up to t and then, subsequently, separately for B and C. However, we continue to use 1951 population densities, which were based on district A boundaries, for B and C. Population densities will not be uniform even within a district and, hence, the values assigned to B and C will be measured with error. However, we always instrument for population density with potential crop yields in our analyses and these statistics are measured at the level of the contemporaneous district, which takes care of the measurement error. We use potential crop yields obtained from the FAO GAEZ database for 42 crops to predict population densities. These yields are provided at a resolution of 0.0174 degrees (1.943km. at the equator) and can be mapped to the Indian district.

For the analysis across regions at the country level (Figure 5) we use gridded population data from the year 2000, which are available at a resolution of 30 sec (1km. at the equator) from the NASA SEDAC Gridded Population of the World version 4. To predict population densities at the country level, we first compute population density statistics from the NASA SEDAC database at the district (second administrative unit) level. We then predict population densities at the district level using potential crop yields obtained from the FAO GAEZ database for the 42 crops. We finally take the population weighted average across all districts to construct a measure of (predicted) population density at the country level.

To construct population densities at the district level for the cross-regional analysis with DHS data, we start with the cluster-level statistics, which are also derived from the NASA SEDAC database. There are 25-30 households in each DHS cluster. We then average across all clusters to construct district-level population density statistics. We finally use potential crop yields obtained from the FAO GAEZ database to predict the population densities at the district level in Figure 6. The potential crop yields are used as statistical instruments in Appendix Table B8.

2. Labor force participation: The NSS labor force participation statistic is derived from the usual activity status of all working-age adults in the household. An individual is coded as participating in the labor force if they work in a household enterprise, are self-employed, work as a regular salaried or casual worker, had worked in the past but do not currently due to sickness or other reasons, and did not seek but are available for work. An individual is coded as not participating if they attend an educational institution, attend domestic duties only, or are otherwise unavailable for work. Individual responses are aggregated up to the district level.

The ILO UN STATS database provides estimates of labor force participation for the rural 15+ population in 2005, separately for men and women. This country-level statistic is used directly in

Figure 5. The DHS provides information on employment and not labor force participation. The DHS survey elicits the following information for each respondent: whether they are currently employed; i.e. worked in the past 7 days, worked in the past 12 months but are not currently employed, or were not employed in the past 12 months. We code an individual as working if they are currently employed or worked in the past 12 months. The individual responses are aggregated up to the district level to construct unemployment rates in Figure 6 and Appendix Table B8.

3. Additional NSS variables:

(a) Vegetarianism: If a household spent a positive amount in the preceding month on the consumption of chicken, pork, beef, goat, or eggs, then the vegetarianism variable is set to zero (one otherwise). We do not include fish in the list of non-vegetarian items because even Brahmins eat fish in coastal regions, where the bulk of this food product is consumed (Srinivas, 1967). The sample is restricted to rural Hindu households who did not have a religious ceremony in the 30 days preceding the survey.

(b) Teetotalism: If a household spent a positive amount in the preceding month on country liquor, foreign liquor, beer or toddy, then the teetotalism variable is set to zero (one otherwise). The sample is restricted to rural Hindu households who did not have a religious ceremony in the 30 days preceding the survey.

(c) Expenditures and prices: For expenditures, we compute the amount spent in the last 30 days on rice, wheat, other cereals and their substitutes, pulses and their derivative products, milk and associated products, edible oils, meat and fish, vegetables, fruits, spices, tobacco, alcohol, fuel and light, clothing including footwear, education, medical services, entertainment, toiletries, transport, rent, and taxes. The NSS uses either a 7 day, 30 day, or yearly recall over different survey rounds and different consumption goods. We do an imputation to convert different reporting periods to a 30 day recall. For the price of meat and alcohol, we compute the consumption-weighted Paasche index, which is calculated as a weighted average of the price of different items, using the expenditure shares of the items as weights. The price for each item is calculated as its value divided by the quantity. For meat, we include goat meat/mutton, eggs, pork, beef/buffalo meat, other meat, and chicken. For alcohol, we include toddy, country liquor, beer, and foreign liquor.

(d) Female wages: The daily wage is recorded for each individual over the past seven days. We take the average over all working days to construct the mean wage.

(e) Female education: The NSS records each individual's education in the following categories: primary school completion, middle school completion, secondary school completion, and college graduate. We convert these categories into years of education, as follows: primary = 4 years, middle = 8 years, secondary = 12 years, and graduate = 16 years.

4. Additional DHS variables:

(a) Marriage and fertility: The district-level marriage rate is constructed as the fraction of women aged 15-49 who are married. The fertility rate is measured by the average number of children ever born and the average number of surviving children for women aged 40+.

(b) Decision-making and autonomy: The DHS survey asks who usually makes health care and expenditure decisions in the household. If the female respondent and her spouse both decide, then the variable is coded as one (zero otherwise). The survey also asks whether the respondent needs permission to visit her relatives. If the answer is negative, then the variable is coded as one (zero otherwise).

5. Additional IHDS variables:

(a) Decision-making and autonomy: The IHDS asks who in the family decides the following: how many children to have, what to cook on a daily basis, what items to buy, and the choice of treatment for sick children. If the female respondent has a say in a given decision, then the variable is coded as one (zero otherwise). The survey also asks whether the respondent needs permission to go out. If the answer is no, then the variable is coded as one (zero otherwise).

(b) Status signals: If a household spent a positive amount in the preceding month on the consumption of meat or eggs, then the vegetarianism variable is coded as zero (one otherwise). If the household spent a positive amount in the preceding month on intoxicants, including alcohol, *pan*, and tobacco, then the teetotalism variable is coded as zero (one otherwise). Note that the IHDS does not provide separate information on alcohol consumption. If the women in the household practice *ghunghat*, *purdah*, or *pallu* then the veiling variable is coded as one (zero otherwise).

6. Rainfall: The rainfall variable that we use for the instrumental variable analysis is constructed using the Climate Research Unit Time Series (CRU TS) gridded precipitation data (Harris et al., 2020), which is available at a resolution of $0.5^\circ \times 0.5^\circ$ each month over the 1901-2018 period. We first calculate total annual rainfall from the monthly data. We then use the spatial district maps to calculate average annual rainfall within each district in each year.

Table A1: DHS Countries and Sample Years

Country	Sample years
<i>Sub-Saharan Africa</i>	
Angola	2015
Burkina Faso	1999, 2003, 2010
Benin	1996, 2001, 2012
Burundi	2010, 2016
Democratic Republic of Congo	2007, 2013
Cote d'Ivoire	1998, 2012
Cameroon	2004, 2011
Ethiopia	2000, 2005, 2011, 2016
Gabon	2012
Ghana	1998, 2003, 2008, 2014
Guinea	1999, 2005, 2012
Kenya	2003, 2008, 2014
Liberia	2007, 2013
Lesotho	2004, 2009, 2014
Mali	1996, 2001, 2006, 2012
Malawi	2000, 2004, 2010, 2015
Mozambique	2011
Nigeria	2003, 2008, 2013
Namibia	2000, 2006, 2013
Rwanda	2005, 2010, 2014
Sierra Leone	2008, 2013
Senegal	2005, 2010, 2012, 2015, 2016
Chad	2014
Tanzania	1999, 2010, 2015
Zambia	2007, 2013
Zimbabwe	1999, 2005, 2010, 2015
<i>South and South East Asia</i>	
Bangladesh	2004, 2007, 2011, 2014
India	2015
Cambodia	2000, 2005, 2010, 2014
Myanmar	2016
Nepal	2001, 2006, 2011, 2016
Philippines	2003, 2008, 2017
Pakistan	2006

B Cross-Sectional Evidence

Robinson Procedure

Consider the following semi-parametric estimating equation:

$$y_j = f(Z_j) + X_j\beta + \epsilon_j$$


where y_j is an outcome, such as FLFP in district j , Z_j is population density, X_j is a vector of covariates, such as state fixed effects, that need to be partialled out prior to nonparametric estimation of the $y_j - Z_j$ association and ϵ_j is a mean-zero disturbance term. The Robinson Robinson (1988) procedure is implemented as follows:

Step 1. Separately regress y_j and each element of the X_j vector nonparametrically on Z_j .

Step 2. Regress the residuals from the first equation, $\hat{\xi}_y$, on the residuals from the other equations, $\hat{\xi}_X$, using a linear specification without a constant term to estimate $\hat{\beta}$.

Step 3. Nonparametrically regress $y_j - (X_j - \bar{X})\hat{\beta}$ on Z_j , where $\bar{X}$ is the sample mean of each element in the vector of covariates.

Figure B1: Population, size of the economy, and size of the medical facilities



Source: ICRISAT district level data, 2011 population census, Village Directory (Asher et al., 2021)

Rural population and agriculture output value is measured in the year 2011.

Number of doctors per facility is top-coded at 30.

To standardize a variable, we subtract its mean and divide by its standard deviation.

Table B1: Health services and facility size (IHDS)

Dep. var.:	procedures (1)	tests (2)	equipment (3)
<i>Panel A: Primary health centers</i>			
facility size	0.534*** (0.097)	0.587*** (0.205)	0.990*** (0.169)
Dep. var. mean	16.090	13.202	18.282
Observations	535	535	535
<i>Panel B: Community health centers</i>			
facility size	0.319*** (0.087)	0.459*** (0.162)	0.243** (0.097)
Dep. var. mean	19.696	17.574	24.549
Observations	204	204	204
<i>Panel C: Rural public hospitals</i>			
facility size	1.305*** (0.080)	1.418*** (0.115)	1.691*** (0.111)
Dep. var. mean	16.137	13.425	18.594
Observations	160	160	160

Source: IHDS Medical Facility Survey, 2011

Health facility size is measured by the number of doctors in place (top-coded at 30).

Procedures include child immunizations, contraceptive services, prenatal care, incision of abscesses and boils, saline IV, setting broken bones, treating gynaecological conditions, treating STDs/STIs, DOTS for tuberculosis, eye exams, treating diarrhea, changing a wound dressing, stitching wounds, treating malaria, treating minor illnesses like fever, rabies injections, childbirth, abortion, blood transfusion, cataract surgery, abdominal surgery, and heart surgery.

Tests include pregnancy, blood pressure, blood sugar, haemoglobin, white blood cell count, HIV/AIDS, cholesterol, urine culture, stool, chlorine level in water, malaria, cerebral malaria, TB, and pap smear.

Equipment includes stethoscope, thermometer, vaginal speculum, sonograph/ultrasound, x-ray machine, blood pressure gauge, oxygen, otoscope for ear exam, ophthalmoscope for eye exam, delivery kit, forceps, partograph for tracking delivery, IV stand, laryngoscope for throat, catheter (urethral), microscope, centrifuge, refrigerator, cold chest, ECG monitor, ambulance, wheelchair, stretcher on a trolley, computer, internet connection, landline telephone, and mobile phone communicating with patients.

The dependent variable is the number of procedures, tests, equipment (based on the IHDS list provided above) in the facility.

Table B2: Rural labor force non-participation, 25-65 age range (Indian districts, NSS)

Dep. variable	rural labor force non-participation					
	female			male		
Caste group:	all	high	low	all	high	low
	(1)	(2)	(3)	(4)	(5)	(6)
Population density	0.119*** (0.027)	0.119*** (0.030)	0.111*** (0.026)	-0.008*** (0.003)	-0.011*** (0.004)	-0.005 (0.005)
Population density $\times$ time trend	-0.003*** (0.001)	-0.003*** (0.001)	-0.002*** (0.001)	0.000*** (0.000)	0.000*** (0.000)	0.000* (0.000)
Kleibergen-Paap Wald F-statistic	21.77	27.92	16.83	21.58	27.93	12.22
Dep. var. mean	0.652	0.670	0.585	0.035	0.035	0.033
Observations	3408	3401	3297	3409	3402	3295

Source: NSS (“thick” and “thin” rounds) and 1951 population census

Population density in 1951, measured in logs, is instrumented using FAO GAEZ potential crop yields.

State and NSS round fixed effects are included in the estimating equation.

Standard errors are clustered at the level of 1981 district boundaries.

Table B3: Rural labor force non-participation, NSS “thick” rounds (Indian districts, NSS))

Dep. variable	rural labor force non-participation					
	female			male		
Caste group:	all	high	low	all	high	low
	(1)	(2)	(3)	(4)	(5)	(6)
Population density	0.135*** (0.030)	0.144*** (0.033)	0.101*** (0.027)	0.001 (0.004)	-0.001 (0.004)	0.001 (0.008)
Population density $\times$ time trend	-0.004*** (0.001)	-0.004*** (0.001)	-0.002*** (0.001)	0.000** (0.000)	0.000 (0.000)	0.000* (0.000)
Kleibergen-Paap Wald F-statistic	26.62	28.78	26.95	26.58	28.70	26.97
Dep. var. mean	0.664	0.697	0.598	0.082	0.089	0.070
Observations	2080	2073	2060	2082	2074	2059

Source: NSS (“thick” rounds) and 1951 population census

Population density in 1951, measured in logs, is instrumented using FAO GAEZ potential crop yields.

State and NSS round fixed effects are included in the estimating equation.

Standard errors are clustered at the level of 1981 district boundaries.

Table B4: Rural labor force non-participation, Muslims and other religions (Indian districts, NSS)

Dep. variable	rural labor force non-participation			
	Muslims		other religions	
	female (1)	male (2)	female (3)	male (4)
Population density	0.087*** (0.022)	0.015* (0.009)	0.130*** (0.047)	-0.015 (0.015)
Population density $\times$ time trend	-0.003*** (0.001)	0.000 (0.000)	-0.003 (0.002)	0.001 (0.001)
Kleibergen-Paap Wald F-statistic	10.79	10.86	13.41	12.48
Dep. var. mean	0.761	0.077	0.629	0.085
Observations	2622	2625	1782	1765

Source: NSS (“thick” and “thin” rounds) and 1951 population census
Population density in 1951, measured in logs, is instrumented using FAO GAEZ potential crop yields.
State and NSS round fixed effects are included in the estimating equation.
Standard errors are clustered at the level of 1981 district boundaries.

Table B5: Rural expenditures and prices (Indian districts, NSS)

Dependent var.	log total expenditures (1)	log food expenditures (2)	log meat price (3)	log alcohol price (4)
Population density	0.085*** (0.031)	0.043** (0.021)	-0.249* (0.131)	-0.046 (0.081)
Population density $\times$ time trend	-0.001 (0.001)	-0.001 (0.001)	0.005 (0.007)	0.003 (0.004)
Kleibergen-Paap Wald F-statistic	12.23	25.33	27.21	22.40
Dep. var. mean	6.818	6.367	1.966	0.814
Observations	1765	2083	2057	1968

Source: NSS (“thick” rounds) and 1951 population census
Population density in 1951, measured in logs, is instrumented using FAO GAEZ potential crop yields.
State and NSS round fixed effects are included in the estimating equation.
Standard errors are clustered at the level of 1981 district boundaries.

Table B6: Rural patriarchal norms (Indian districts, IHDS)

Dep. variable	how many children (1)	whom children marry (2)	what to cook (3)	what to buy (4)	treatment of sick children (5)	does not need permission to go out (6)
Population density	0.039 (0.025)	0.000 (0.016)	0.028 (0.038)	0.026* (0.016)	0.019 (0.025)	0.040 (0.026)
Kleibergen-Paap Wald F-statistic	7.20	7.20	7.20	7.20	7.20	7.20
Dep. var. mean	0.197	0.111	0.688	0.106	0.259	0.201
Observations	237	237	237	237	237	237

Source: IHDS and 1951 population census

Patriarchal norms are measured by the fraction of women who report having a say in household decisions and not needing permission to go out.

Population density in 1951, measured in logs, is instrumented by FAO GAEZ potential crop yields.

State fixed effects are included in the estimating equation.

Table B7: Rural status signaling (Indian districts, IHDS)

Dep. var.:	FLFNP (1)	vegetarianism (2)	teetotalism (3)	veiling (4)
Population density	0.064** (0.030)	0.037 (0.040)	0.050** (0.021)	0.110** (0.051)
Kleibergen-Paap Wald F-statistic	7.20	7.20	7.20	7.20
Dep. var. mean	0.679	0.687	0.306	0.612
Observations	237	237	237	237

Source: IHDS and 1951 population census

Population density in 1951, measured in logs, is instrumented by FAO GAEZ potential crop yields.

State fixed effects are included in the estimating equation.

Table B8: Rural unemployment across regions (district data, DHS)

Dep. variable	rural unemployment							
	female				male			
	Africa		Asia		Africa		Asia	
	All (1)	all (2)	only India (3)	excluding India (4)	all (5)	all (6)	only India (7)	excluding India (8)
Population density	-0.012 (0.008)	0.042*** (0.005)	0.028*** (0.009)	0.038*** (0.005)	-0.000 (0.002)	0.001 (0.001)	0.003 (0.002)	-0.001 (0.001)
Kleibergen-Paap Wald F-statistic	22.02	17.55	10.73	21.48	22.41	17.60	10.73	21.58
Dep. var. mean	0.342	0.377	0.658	0.304	0.026	0.018	0.037	0.013
Observations	5943	2801	579	2222	6328	2740	579	2161

Source: DHS and NASA SEDAC

Population density in 2000, measured in logs, is instrumented by FAO GAEZ potential crop yields.

First administrative unit (state) fixed effects and survey year effects are included in the estimating equation.

C The Model

C.1 Proof of Proposition 1

Using the same notation as in Section 3, household i in group $k \in \{H, L\}$ chooses $c_{i,k}$ to maximize

$$\log(y_{i,k} - c_{i,k}) + \frac{\mathbb{C}_k}{\mathbb{C}_k + \mathbb{C}_{-k}} \cdot 2v.$$

Since all N households in each group have the same income endowment, this is a symmetric equilibrium and, hence, the optimal choice of $c_{i,k}$ is determined by the following first-order condition:

$$\frac{1}{y_k - c_k} = \frac{c_{-k}}{(c_k + c_{-k})^2} \cdot 2 \frac{v}{N}.$$

This constitutes a system of two equations with two unknowns: c_H, c_L . To solve these equations, divide one by the other and collect terms to obtain:

$$\frac{c_k}{y_k} = \frac{c_{-k}}{y_{-k}}$$

and then substitute back in the first-order condition to derive an equation with a single unknown, c_k :

$$\frac{1}{y_k - c_k} = \frac{y_k y_{-k}}{c_k (y_k + y_{-k})^2} \cdot 2 \frac{v}{N}.$$

Denote $K \equiv \frac{(y_H + y_L)^2}{y_H y_L}$, $w \equiv \frac{N}{2v}$. The preceding equation can then be rewritten as

$$\frac{1}{y_k - c_k} = \frac{1}{c_k K w},$$

which implies that

$$c_k = \frac{y_k}{1 + K w}.$$

Taking the average over $k = H, L$:

$$\bar{c} = \frac{\bar{y}}{1 + K w}.$$

Since w is decreasing in v , it follows immediately that $\bar{c}$ is increasing in v . To derive the corresponding implications with respect to $\bar{y} = \frac{y_H + y_L}{2}$ and $\Delta y = \frac{y_H - y_L}{2}$, we rewrite K as a function of $\bar{y}, \Delta y$:

$$K \equiv \frac{(y_H + y_L)^2}{y_H y_L} = \frac{4\bar{y}^2}{\bar{y}^2 - \Delta y^2}.$$

Differentiating K with respect to Δy and $\bar{y}$:

$$\frac{\partial K}{\partial \Delta y} = \frac{8\bar{y}^2 \Delta y}{(\bar{y}^2 - \Delta y^2)^2} > 0$$

$$\frac{\partial K}{\partial \bar{y}} = \frac{-8\bar{y} \Delta y^2}{(\bar{y}^2 - \Delta y^2)^2} < 0$$

Since K is increasing in Δy , it follows immediately that $\bar{c}$ is decreasing in Δy . It is also straightforward to verify that $\bar{c}$ is increasing in $\bar{y}$ because K is decreasing in $\bar{y}$ and $\bar{y}$ appears in the numerator of the $\bar{c}$

expression. This completes the proof of Proposition 1.

To compare the magnitude of the different partial effects, we derive the following expressions:

$$\begin{aligned}\frac{\partial \bar{c}}{\partial \bar{y}} &= \frac{1}{1 + Kw} - \frac{\bar{y}}{(1 + Kw)^2} \cdot w \frac{\partial K}{\partial \bar{y}} \\ \frac{\partial c_L}{\partial \bar{y}} &= \frac{1}{1 + Kw} - \frac{\bar{y} - \Delta y}{(1 + Kw)^2} \cdot w \frac{\partial K}{\partial \bar{y}} \\ \frac{\partial c_H}{\partial \bar{y}} &= \frac{1}{1 + Kw} - \frac{\bar{y} + \Delta y}{(1 + Kw)^2} \cdot w \frac{\partial K}{\partial \bar{y}} \\[10pt] \frac{\partial \bar{c}}{\partial \Delta y} &= \frac{-\bar{y}}{(1 + Kw)^2} \cdot w \frac{\partial K}{\partial \Delta y} \\ \frac{\partial c_L}{\partial \Delta y} &= \frac{-(\bar{y} - \Delta y)}{(1 + Kw)^2} \cdot w \frac{\partial K}{\partial \Delta y} - \frac{1}{1 + Kw} \\ \frac{\partial c_H}{\partial \Delta y} &= \frac{-(\bar{y} + \Delta y)}{(1 + Kw)^2} \cdot w \frac{\partial K}{\partial \Delta y} + \frac{1}{1 + Kw}\end{aligned}$$

If we assume that $\frac{\Delta y}{(1 + Kw)^2} \approx 0$, then it follows that:

$$\frac{\partial \bar{c}}{\partial \bar{y}} = \frac{\partial c_L}{\partial \bar{y}} = \frac{\partial c_H}{\partial \bar{y}}$$

$$\left| \frac{\partial c_L}{\partial \Delta y} \right| > \left| \frac{\partial \bar{c}}{\partial \Delta y} \right| > \left| \frac{\partial c_H}{\partial \Delta y} \right|$$

Figure C1: Wages against rainfall, within caste-gender-occupation cells




Table C1: First stage regression

Dep. variable	$\bar{y}$ (1)	Δy (2)
$\bar{y}_{IV}$	1.075*** (0.073)	0.243** (0.115)
Δy_{IV}	0.038 (0.052)	0.809*** (0.076)
F -statistic excluded instruments	127.20 [0.000]	89.63 [0.000]
Observations	2832	2832

Source: NSS ("thick" and "thin" rounds)

District and NSS round fixed effects are included in the estimating equation.
Standard errors are clustered at the level of 1981 district boundaries.

Table C2: Female labor force non-participation within districts over time (“thick” rounds only)

Dep. variable	FLFNP					
	OLS			IV		
	all	high	low	all	high	low
Regression:	(1)	(2)	(3)	(4)	(5)	(6)
mean potential income	0.261*** (0.081)	0.217** (0.084)	0.336*** (0.102)	0.526*** (0.177)	0.457** (0.182)	0.869*** (0.263)
caste-gap in potential income	-0.082 (0.052)	-0.001 (0.054)	-0.209*** (0.067)	-0.288* (0.153)	-0.066 (0.160)	-0.626*** (0.220)
Kleibergen-Paap Wald F-statistic	—	—	—	50.46	50.46	50.46
Kleibergen-Paap LM statistic	—	—	—	63.83	63.83	63.83
Dep. var. mean	0.657	0.691	0.589	0.657	0.691	0.589
Observations	1522	1522	1522	1522	1522	1522

Source: NSS (“thick” rounds)

District and NSS round fixed effects are included in the estimating equation.

Standard errors are clustered at the level of 1981 district boundaries.

Table C3: Female labor force non-participation within districts over time (national-level population shares)

Dep. variable	FLFNP					
	OLS			IV		
	all	high	low	all	high	low
Regression:	(1)	(2)	(3)	(4)	(5)	(6)
mean potential income	0.232*** (0.058)	0.167*** (0.060)	0.342*** (0.072)	0.785*** (0.157)	0.803*** (0.175)	0.977*** (0.195)
caste-gap in potential income	-0.095*** (0.036)	-0.019 (0.039)	-0.216*** (0.047)	-0.204* (0.117)	-0.058 (0.115)	-0.553*** (0.173)
Kleibergen-Paap Wald F-statistic	—	—	—	77.03	77.03	77.03
Kleibergen-Paap LM statistic	—	—	—	80.97	80.97	80.97
Dep. var. mean	0.650	0.685	0.585	0.650	0.685	0.585
Observations	2894	2894	2894	2894	2894	2894

Source: NSS (“thick” and “thin” rounds)

The instruments are constructed using national-level population shares.

District and NSS round fixed effects are included in the estimating equation.

Standard errors are clustered at the level of 1981 district boundaries.

Table C4: Female labor force non-participation within districts over time (population density interacted with time period effects)

Dep. variable	FLFNP					
	OLS			IV		
	all	high	low	all	high	low
Regression:	(1)	(2)	(3)	(4)	(5)	(6)
Caste group:						
mean potential income	0.266*** (0.057)	0.197*** (0.059)	0.367*** (0.073)	0.870*** (0.158)	0.859*** (0.175)	1.024*** (0.198)
caste-gap in potential income	-0.108*** (0.036)	-0.032 (0.039)	-0.216*** (0.048)	-0.283** (0.120)	-0.142 (0.122)	-0.553*** (0.180)
Kleibergen-Paap Wald F-statistic	—	—	—	81.66	81.66	81.66
Kleibergen-Paap LM statistic	—	—	—	82.75	82.75	82.75
Dep. var. mean	0.648	0.683	0.582	0.648	0.683	0.582
Observations	2831	2831	2831	2831	2831	2831

Source: NSS (“thick” and “thin” rounds)

Population density interacted with time period effects is included in the estimating equation.

District and NSS round fixed effects are included in the estimating equation.

Standard errors are clustered at the level of 1981 district boundaries.

Table C5: Female labor force non-participation within districts over time (accounting for income from land)

Dep. variable	FLFNP					
	OLS			IV		
	all	high	low	all	high	low
Regression:	(1)	(2)	(3)	(4)	(5)	(6)
Caste group:						
mean potential income	0.244*** (0.059)	0.175*** (0.060)	0.349*** (0.075)	0.891*** (0.157)	0.872*** (0.174)	1.049*** (0.197)
caste-gap in potential income	-0.098*** (0.038)	-0.021 (0.040)	-0.213*** (0.048)	-0.278** (0.122)	-0.139 (0.126)	-0.560*** (0.180)
Kleibergen-Paap Wald F-statistic	—	—	—	85.63	85.63	85.63
Kleibergen-Paap LM statistic	—	—	—	83.18	83.18	83.18
Dep. var. mean	0.648	0.683	0.582	0.648	0.683	0.582
Observations	2832	2832	2832	2832	2832	2832

Source: NSS (“thick” and “thin” rounds)

Caste-specific land incomes are included in the estimating equation.

District and NSS round fixed effects are included in the estimating equation.

Standard errors are clustered at the level of 1981 district boundaries.

D Structural Estimation and Policy Simulations

Benchmark model setup: Household i belonging to caste group k chooses $\tau_{i,sk}$, $\tau_{i,uk}$, $\xi_{i,kg}$ to maximize:

$$\log(y_{i,k} - c_{i,sk} - c_{i,uk}) + \frac{\mathbb{C}_k}{\mathbb{C}_k + \mathbb{C}_{-k}} \cdot 2v,$$

subject to the household's potential income

$$y_{i,k} = \sum_g w_{sg} \xi_{i,kg} + w_{ug}(1 - \xi_{i,kg}) - e_{kg}(\xi_{i,kg}),$$

with the signaling costs expressed as follows:

$$\begin{aligned} c_{i,sk} &= w_{skf} \xi_{i,kf} \tau_{i,sk} \eta_{sk} \\ c_{i,uk} &= w_{ukf} (1 - \xi_{i,kf}) \tau_{i,uk} \eta_{uk} \\ c_{i,k} &= \left(c_{i,sk}^\phi + c_{i,uk}^\phi \right)^{\frac{1}{\phi}} \\ \mathbb{C}_k &= \sum_{j=1}^{N_k} c_{jk} \\ e_{kg}(\xi_{i,kg}) &= \beta_{kg} \xi_{i,kg}^{\nu_{kg}} \end{aligned}$$

Notice that the preceding specification allows the skilled and unskilled signaling costs, $c_{i,sk}$ and $c_{i,uk}$, to enter as imperfect substitutes or even complements in the status function. This assumption is needed to derive independent equations for $\tau_{i,sk}$, $\tau_{i,uk}$, as seen below. We will not need to estimate ϕ and thus this parameter could, in principle, be arbitrarily close to one, in which case $c_{i,sk}$, $c_{i,uk}$ would be (close to) perfect substitutes.

First-order conditions:

FOC $_{\tau_{i,sk}}$:

$$-\frac{1}{y_{i,k} - c_{i,sk} - c_{i,uk}} \frac{\partial c_{i,sk}}{\partial \tau_{i,sk}} + 2v \cdot \frac{\mathbb{C}_{-k}}{(\mathbb{C}_k + \mathbb{C}_{-k})^2} \frac{\partial \mathbb{C}_k}{\partial \tau_{i,sk}} = 0$$

FOC $_{\tau_{i,uk}}$:

$$-\frac{1}{y_{i,k} - c_{i,sk} - c_{i,uk}} \frac{\partial c_{i,uk}}{\partial \tau_{i,uk}} + 2v \cdot \frac{\mathbb{C}_{-k}}{(\mathbb{C}_k + \mathbb{C}_{-k})^2} \frac{\partial \mathbb{C}_k}{\partial \tau_{i,uk}} = 0$$

FOC $_{\xi_{i,kf}}$:

$$\frac{1}{y_{i,k} - c_{i,sk} - c_{i,uk}} \left[\frac{\partial y_{i,k}}{\partial \xi_{i,kf}} - \frac{\partial c_{i,sk}}{\partial \xi_{i,kf}} - \frac{\partial c_{i,uk}}{\partial \xi_{i,kf}} \right] + 2v \cdot \frac{\mathbb{C}_{-k}}{(\mathbb{C}_k + \mathbb{C}_{-k})^2} \frac{\partial \mathbb{C}_k}{\partial \xi_{i,kf}} = 0$$

FOC $_{\xi_{i,km}}$:

$$\frac{1}{y_{i,k} - c_{i,sk} - c_{i,uk}} \frac{\partial y_{i,k}}{\partial \xi_{i,km}} = 0$$

Parameterizing the cost of education function by $e_{kg}(\xi_{i,kg}) = \beta_{kg} \xi_{i,kg}^{\nu_{kg}}$, and solving the partial derivatives, the first-order conditions can be written as

FOC $_{\tau_{i,sk}}$:

$$-\frac{w_{skf} \xi_{i,kf} \eta_{sk}}{y_{i,k} - c_{i,sk} - c_{i,uk}} + 2v \cdot \frac{\mathbb{C}_{-k}}{(\mathbb{C}_k + \mathbb{C}_{-k})^2} c_{i,k}^{1-\phi} c_{i,sk}^{\phi-1} w_{skf} \xi_{i,kf} \eta_{sk} = 0$$

FOC $_{\tau_{i,uk}}$:

$$-\frac{w_{ukf} (1 - \xi_{i,kf}) \eta_{uk}}{y_{i,k} - c_{i,sk} - c_{i,uk}} + 2v \cdot \frac{\mathbb{C}_{-k}}{(\mathbb{C}_k + \mathbb{C}_{-k})^2} c_{i,k}^{1-\phi} c_{i,uk}^{\phi-1} w_{ukf} (1 - \xi_{i,kf}) \eta_{uk} = 0$$

FOC $_{\xi_{i,kf}}$:

$$\begin{aligned} & \frac{1}{y_{i,k} - c_{i,sk} - c_{i,uk}} \left[w_{skf} (1 - \tau_{i,ks} \eta_{sk}) - w_{ukf} (1 - \tau_{i,uk} \eta_{uk}) - \beta_{kf} \nu_{kf} \xi_{i,kf}^{\nu_{kf}-1} \right] \\ & + 2v \cdot \frac{\mathbb{C}_{-k}}{(\mathbb{C}_k + \mathbb{C}_{-k})^2} c_{i,k}^{1-\phi} [c_{i,sk}^{\phi-1} w_{skf} \tau_{i,sk} - c_{i,uk}^{\phi-1} w_{ukf} \tau_{i,uk} \eta_{uk}] = 0 \end{aligned}$$

FOC $_{\xi_{i,km}}$:

$$\frac{1}{y_{i,k} - c_{i,sk} - c_{i,uk}} [w_{skm} - w_{ukm} - \beta_{km} \nu_{km} \xi_{i,km}^{\nu_{km}-1}] = 0$$

Notice that the first two first-order conditions would collapse to a single equation if we set ϕ equal to one. With this more general specification, we have two distinct equations that will allow us to solve for $\tau_{i,sk}$, $\tau_{i,uk}$. Notice also that we can pin down male education from FOC $_{\xi_{i,km}}$:

$$\xi_{i,km} = \left[\frac{w_{skm} - w_{ukm}}{\beta_{km} \nu_{km}} \right]^{\frac{1}{\nu_{km}-1}} \quad (\text{D.1})$$

Male education in the preceding equation depends only on caste-skill level wages and model parameters that

do not vary across households. Hence, we can replace $\xi_{i,km}$ with ξ_{km} on the left hand side of equation (D.1).

Since there is no heterogeneity within castes in a given district-time period, we can also set $\xi_{i,kf} = \xi_{kf}$, $\tau_{i,sk} = \tau_{sk}$, and $\tau_{i,uk} = \tau_{uk}$. It then follows that $y_{i,k} = y_k$, $c_{i,sk} = c_{sk}$, $c_{i,uk} = c_{uk}$, $c_{i,k} = c_k$, and $\mathbb{C}_k = N_k c_k$, which allows us to rewrite the first-order conditions in terms of τ_{sk} , τ_{uk} , and ξ_{kf} :

$$\begin{aligned}
\text{FOC}_{\tau_{sk}} : & \\
& -\frac{w_{skf}\xi_{kf}\eta_{sk}}{y_k - c_{sk} - c_{uk}} + 2v \cdot \frac{N_{-k}c_{-k}}{(N_k c_k + N_{-k}c_{-k})^2} c_k^{1-\phi} c_{sk}^{\phi-1} w_{skf}\xi_{kf}\eta_{sk} = 0 \\
\text{FOC}_{\tau_{uk}} : & \\
& -\frac{w_{ukf}(1 - \xi_{kf})\eta_{uk}}{y_k - c_{sk} - c_{uk}} + 2v \cdot \frac{N_{-k}c_{-k}}{(N_k c_k + N_{-k}c_{-k})^2} c_k^{1-\phi} c_{uk}^{\phi-1} w_{ukf}(1 - \xi_{kf})\eta_{uk} = 0 \\
\text{FOC}_{\xi_{kf}} : & \\
& \frac{1}{y_k - c_{sk} - c_{uk}} \left[w_{skf}(1 - \tau_{sk}\eta_{sk}) - w_{ukf}(1 - \tau_{uk}\eta_{uk}) - \beta_{kf}\nu_{kf}\xi_{kf}^{\nu_{kf}-1} \right] \\
& + 2v \cdot \frac{N_{-k}c_{-k}}{(N_k c_k + N_{-k}c_{-k})^2} c_k^{1-\phi} [c_{sk}^{\phi-1} w_{skf}\tau_{sk}\eta_{sk} - c_{uk}^{\phi-1} w_{ukf}\tau_{uk}\eta_{uk}] = 0
\end{aligned}$$

By inspection of $\text{FOC}_{\tau_{sk}}$ and $\text{FOC}_{\tau_{uk}}$, $c_{sk} = c_{uk}$, which can be rewritten as:

$$\frac{\xi_{kf}\tau_{sk}}{(1 - \xi_{kf})\tau_{uk}} = \frac{w_{ukf}\eta_{uk}}{w_{skf}\eta_{sk}} \quad (\text{D.2})$$

This also implies that we can drop one first-order condition, say $\text{FOC}_{\tau_{uk}}$, and retain $\text{FOC}_{\tau_{sk}}$, replacing c_{uk} with c_{sk} :

$$\frac{1}{y_k - 2c_{sk}} = 2v \cdot \frac{N_{-k}c_{-k}}{(N_k c_k + N_{-k}c_{-k})^2} c_k^{1-\phi} c_{sk}^{\phi-1} \quad (\text{D.3})$$

Using the definition $c_k = (c_{sk}^\phi + c_{uk}^\phi)^{\frac{1}{\phi}}$, and substituting $c_{sk} = c_{uk}$, gives $c_k = 2^{\frac{1}{\phi}} c_{sk}$. Hence, we can rewrite (D.3):

$$\frac{1}{y_k - 2c_{sk}} = v \frac{N_{-k}c_{s,-k}}{(N_k c_{sk} + N_{-k}c_{s,-k})^2} \quad (\text{D.4})$$

The first-order condition $\text{FOC}_{\tau_{sk}}$ can be simplified to (D.4), which is independent of ϕ . Similarly, noting that $c_{sk} = c_{uk}$ and $c_k = 2^{\frac{1}{\phi}} c_{sk}$, and using (D.4), we can rewrite $\text{FOC}_{\xi_{kf}}$:

$$w_{skf}(1 - \tau_{sk}\eta_{sk}) - w_{ukf}(1 - \tau_{uk}\eta_{uk}) - \beta_{kf}\nu_{kf}\xi_{kf}^{\nu_{kf}-1} + w_{skf}\tau_{sk}\eta_{sk} - w_{ukf}\tau_{uk}\eta_{uk} = 0$$

which gives us an expression for female education that is analogous to (D.1):

$$\xi_{kf} = \left[\frac{w_{skf} - w_{ukf}}{\beta_{kf}\nu_{kf}} \right]^{\frac{1}{\nu_{kf}-1}} \quad (\text{D.5})$$

(D.4) is derived in terms of low-caste households and high-caste households in the village: N_L , N_H .

Although these statistics are unavailable, we do know the low-caste population share in the district, x_L in each time period. Under the maintained assumption that villages are homogeneous in each district-time period, $x_L = \frac{N_L}{N_H + N_L}$ and $f(x_L) \equiv \frac{N_H}{N_L} = \frac{1-x_L}{x_L}$. Focusing on the low-caste group to begin with, the payoff from status for this group can be written as

$$\begin{aligned} v \cdot \frac{N_H c_{sH}}{(N_L c_{sL} + N_H c_{sH})^2} &= \frac{v}{N_L} \cdot \frac{\frac{N_H}{N_L} c_{sH}}{(c_{sL} + \frac{N_H}{N_L} c_{sH})^2} \\ &= \tilde{v} \frac{f(x_L) c_{sH}}{(c_{sL} + f(x_L) c_{sH})^2} \end{aligned}$$

The payoff from status for the high-caste group can be derived in the same way, allowing us to rewrite (D.4) for each caste group as

$$\frac{1}{y_L - 2c_{sL}} = \tilde{v} \frac{f(x_L) c_{sH}}{(c_{sL} + f(x_L) c_{sH})^2} \quad (\text{D.6})$$

$$\frac{1}{y_H - 2c_{sH}} = \tilde{v} \frac{c_{sL}}{(c_{sL} + f(x_L) c_{sH})^2} \quad (\text{D.7})$$

Recall from the model set up that y_L, y_H, c_{sL} , and c_{sH} are functions of wages and the endogenous variables, $\tau_{sk}, \tau_{uk}, \xi_{kg}$. Given wages, we thus have four equations for each group k , corresponding to the four endogenous variables: (D.1), (D.2), (D.5), and (D.6)-(D.7).

Wages: We specify a linear aggregate production function $Y = AE$, where A is total factor productivity, and E is aggregate labor. Labor is heterogeneous across three dimensions: (1) gender, (2) caste group, and (3) skill. We use a nested-CES structure for labor aggregation:

$$\begin{aligned} E &= [\theta_f E_f^\rho + \theta_m E_m^\rho]^{\frac{1}{\rho}} \\ E_g &= [\theta_{Lg} E_{Lg}^{\rho_g} + \theta_{Hg} E_{Hg}^{\rho_g}]^{\frac{1}{\rho_g}}, \quad g = \{f, m\} \end{aligned}$$

$$\begin{aligned} E_{kg} &= [\theta_{skg} E_{skg}^{\rho_{kg}} + \theta_{ukg} E_{ukg}^{\rho_{kg}}]^{\frac{1}{\rho_{kg}}}, \quad k = \{H, L\} \\ E_{skf} &= \int_0^{x_k} \xi_{i,kf} (1 - \tau_{i,sk}) di = \xi_{kf} (1 - \tau_{sk}) x_k \\ E_{ukf} &= \int_0^{x_k} (1 - \xi_{i,kf}) (1 - \tau_{i,ku}) di = (1 - \xi_{kf}) (1 - \tau_{ku}) x_k \\ E_{skm} &= \int_0^{x_k} \xi_{i,km} di = \xi_{km} x_k \\ E_{ukm} &= \int_0^{x_k} (1 - \xi_{i,mk}) = (1 - \xi_{km}) x_k \end{aligned}$$

There are eight different wages $W_{skg}, W_{ukg}, g = \{m, f\}, k = \{H, L\}$, which are determined by the following equations:

$$w_{skg} = \frac{\partial Y}{\partial E_{skg}} = \frac{\partial Y}{\partial E} \times \frac{\partial E}{\partial E_g} \times \frac{\partial E_g}{\partial E_{kg}} \times \frac{\partial E_{kg}}{\partial E_{skg}},$$

$$w_{ukg} = \frac{\partial Y}{\partial E_{skg}} = \frac{\partial Y}{\partial E} \times \frac{\partial E}{\partial E_g} \times \frac{\partial E_g}{\partial E_{kg}} \times \frac{\partial E_{kg}}{\partial E_{ukg}}.$$

Solving the partial derivatives:

$$w_{skg} = A\theta_g\theta_{kg}\theta_{skg}E^{1-\rho}E_k^{\rho-\rho_g}E_{kg}^{\rho_g-\rho_{kg}}E_{skg}^{\rho_{kg}-1} \quad (D.8)$$

$$w_{ukg} = A\theta_g\theta_{kg}\theta_{ukg}E^{1-\rho}E_k^{\rho-\rho_g}E_{kg}^{\rho_g-\rho_{kg}}E_{ukg}^{\rho_{kg}-1} \quad (D.9)$$

Estimation: The model is estimated separately in each of the five NSS thick rounds. In each survey round, we construct a (predicted) log population density grid, such that each grid interval contains an equal number of districts. The number of intervals is set equal to 10. In each interval, we compute (i) mean FLFNP, by skill and caste, (ii) mean education, by skill, caste, and gender, and (iii) mean wages, by skill, caste, and gender. We estimate the structural parameters of the model by solving the following problem

$$\min_{\Theta} e(\Theta)'e(\Theta) \quad (D.10)$$

where $e(\Theta)$ is an error vector, computed as the percentage difference between the data moments and the model moments, and Θ is the set of structural parameters.

We assume the following functional form for A and $\tilde{v}$

$$A(p) = \alpha_{A1}p^{\alpha_{A2}}$$

$$\tilde{v}(p) = \alpha_{\tilde{v}1}p^{\alpha_{\tilde{v}2}}$$

where p is the (predicted) log population density in a given interval. Table D1 lists the parameters to be estimated. There are 37 parameters, which we divide in two groups: Group 1 parameters are associated with the cost of education, the non-pecuniary constraints to female labor force participation, and the status function. Group 2 parameters are associated with the aggregate production function. Given the large set

Table D1: Parameters to estimate

Group 1		Group 2	
Social status	$\alpha_{\tilde{v}1}, \alpha_{\tilde{v}2}$	Total factor productivity	α_{A1}, α_{A2}
Cost of education, level	$\beta_{Lf}, \beta_{Hf}, \beta_{Lm}, \beta_{Hm}$	Gender, productivity	θ_f, θ_m
Cost of education, curvature	$\nu_{Lf}, \nu_{Hf}, \nu_{Lm}, \nu_{Hm}$	Gender, elasticity of substitution	ρ
Non-pecuniary constraints	$\eta_{sL}, \eta_{sH}, \eta_{uL}, \eta_{uH}$	Caste-gender productivity	$\theta_{Lf}, \theta_{Lm}, \theta_{Hf}, \theta_{Hm}$
		Caste-gender elasticity of substitution	ρ_f, ρ_m
		Skill-caste-gender productivity	$\theta_{sLf}, \theta_{sLm}, \theta_{sHf}, \theta_{sHm},$
			$\theta_{uLf}, \theta_{uLm}, \theta_{uHf}, \theta_{uHm}$
		Skill-caste-gender elasticity of substitution	$\rho_{Lf}, \rho_{Lm}, \rho_{Hf}, \rho_{Hm}$

of parameters and the objective to find the global minimum, the estimation proceeds in the following steps:

Step 1 Given observed wages and a particular choice of group 1 parameters, solve equations (D.1), (D.2), (D.5), and (D.6)-(D.7) to derive $\tau_{sk}, \tau_{uk}, \xi_{kg}$ in each population density interval. Perform a global search over all parameter values to find the set of parameters that minimizes the distance between the data moments and the model moments with respect to FLFNP and education.

Step 2 Using observed education and FLFNP, which gives us the labor input by caste, gender, and skill, predict wages in each population density interval from equations (D.8) and (D.9), for a particular choice of group 2 parameters. Perform a global search over all parameter values to find the set of parameters that minimizes the distance between observed and predicted wages.

Step 3 Using the iterative algorithm described in Section 5.1, solve for FLFNP, education, and wages for a particular choice of the model's parameters. Using the values obtained in Step 1 and Step 2 as the initial set of parameters, implement a local search procedure to find the set of parameters that minimizes the distance between the data and model moments with respect to FLFNP, education, and wages.

Step 4 Using the parameters estimated in **Step 3** as the initial set of parameters, repeat the local search. Continue until the minimized error from one iteration to the next is below a prespecified threshold value.

For **Step 1** and **Step 2** we use the Differential Evolution algorithm for global optimization (Ardia et al., 2011). For **Step 3** we use Nelder Mead local optimization.

Augmented model with Land income Household i 's potential income is now expressed as

$$y_{i,k} = \sum_g w_{skg} \xi_{i,k,g} + w_{ukg} (1 - \xi_{i,k,g}) + y_{i,k}^d - e_{kg}(\xi_{i,k,g}),$$

where $y_{i,k}^d$ is the income from land.

We assume that output in each district-time period is determined by a Cobb-Douglas aggregate production function: $Y = AD^\gamma E^{1-\gamma}$, where A is total factor productivity, D is aggregate land, and E is aggregate labor. As with the benchmark model, the wage for each skill-caste-gender category is determined by the associated marginal productivity of labor:

$$w_{skg} = \frac{\partial Y}{\partial E_{skg}} = \frac{\partial Y}{\partial E} \times \frac{\partial E}{\partial E_g} \times \frac{\partial E_g}{\partial E_{kg}} \times \frac{\partial E_{kg}}{\partial E_{skg}},$$

$$w_{ukg} = \frac{\partial Y}{\partial E_{ukg}} = \frac{\partial Y}{\partial E} \times \frac{\partial E}{\partial E_g} \times \frac{\partial E_g}{\partial E_{kg}} \times \frac{\partial E_{kg}}{\partial E_{ukg}}.$$

where $\frac{\partial Y}{\partial E_{skg}} = A(1 - \gamma)D^\gamma E^{-\gamma}$.

We use a nested-CES structure, as with labor, to aggregate the different components of land. In the first augmented model, we assume that the representative household in each caste has a different endowment of land. In the second, more flexible, model, we assume that castes have different endowments of irrigated and unirrigated land. We derive the shadow rental rate of land in each case below, which, in turn, allows us to derive the corresponding land income.

Castes with different endowments of land: The nested-CES structure for aggregating the different components of land is expressed as

$$D = (\chi_L D_L^{\rho_D} + \chi_H D_H^{\rho_D})^{\frac{1}{\rho_D}}$$

$$D_L = x_L d_L, \quad D_H = x_H d_H$$

where D_k is the total land owned by group $k = \{H, L\}$. Corresponding to two types of lands, there are two shadow rental rates, R_L, R_H given as

$$\begin{aligned} R_L &= A\gamma\chi_L \left(\frac{E}{D}\right)^{1-\gamma} D^{1-\rho_D} D_L^{\rho_D-1} \\ R_H &= A\gamma\chi_H \left(\frac{E}{D}\right)^{1-\gamma} D^{1-\rho_D} D_H^{\rho_D-1} \end{aligned}$$

The income from land for household i in group k is then given by

$$y_{i,k}^d = R_k d_k,$$

where d_k is the land owned by the representative household in group k . This variable is constructed as the total land owned by caste k , available in the NSS Land and Livestock Holding Survey in 2003, divided by the number of households in that caste, obtained from the Employment and Unemployment Survey in that year.

Castes with different endowments of irrigated and unirrigated land: The nested-CES structure for aggregating the different components of land is now expressed as

$$\begin{aligned} D &= (\chi_L D_L^{\rho_D} + \chi_H D_H^{\rho_D})^{\frac{1}{\rho_D}} \\ D_L &= (\chi_{r,L} D_{r,L}^{\rho_{D,L}} + \chi_{n,L} D_{n,L}^{\rho_{D,L}})^{\frac{1}{\rho_{D,L}}} \\ D_H &= (\chi_{r,H} D_{r,H}^{\rho_{D,H}} + \chi_{n,H} D_{n,H}^{\rho_{D,H}})^{\frac{1}{\rho_{D,H}}} \\ D_{r,L} &= x_L d_{r,L}, \quad D_{n,L} = x_L d_{n,L} \\ D_{r,H} &= x_H d_{r,H}, \quad D_{n,H} = x_H d_{n,H} \end{aligned}$$

where $D_k, D_{r,k}, D_{n,k}$ is total, irrigated, and unirrigated land owned by group $k = \{H, L\}$. Corresponding to four types of lands, there are four shadow rental rates, $R_{r,L}, R_{n,L}, R_{r,H}, R_{n,H}$ given as

$$\begin{aligned} R_{r,L} &= A\gamma\chi_L\chi_{r,L} \left(\frac{E}{D}\right)^{1-\gamma} D^{1-\rho_D} D_L^{\rho_D-\rho_{D,L}} D_{r,L}^{\rho_{D,L}-1} \\ R_{n,L} &= A\gamma\chi_L\chi_{n,L} \left(\frac{E}{D}\right)^{1-\gamma} D^{1-\rho_D} D_L^{\rho_D-\rho_{D,L}} D_{n,L}^{\rho_{D,L}-1} \\ R_{r,H} &= A\gamma\chi_H\chi_{r,H} \left(\frac{E}{D}\right)^{1-\gamma} D^{1-\rho_D} D_H^{\rho_D-\rho_{D,H}} D_{r,H}^{\rho_{D,H}-1} \\ R_{n,H} &= A\gamma\chi_H\chi_{n,H} \left(\frac{E}{D}\right)^{1-\gamma} D^{1-\rho_D} D_H^{\rho_D-\rho_{D,H}} D_{n,H}^{\rho_{D,H}-1} \end{aligned}$$

The income from land for the representative household i in group k is then given by

$$y_{i,k}^d = R_{r,k} d_{r,k} + R_{n,k} d_{n,k},$$

where $d_{r,k}$ and $d_{n,k}$ denote the amount of irrigated and unirrigated land that it owns. As above, these statistics are constructed using land ownership data from the NSS Land and Livestock Holding Survey and

information on the number of households, by caste, from the Employment and Unemployment Survey.

The set of parameters to estimate in the augmented models are given as

Table D2: Parameters to estimate in augmented models

Group 1		Group 2	
Social status	$\alpha_{\tilde{v}_1}, \alpha_{\tilde{v}_2}$	Total factor productivity	$\alpha_{A_1}, \alpha_{A_2}$
Cost of education, level	$\beta_{Lf}, \beta_{Hf}, \beta_{Lm}, \beta_{Hm}$	Gender, productivity	θ_f, θ_m
Cost of education, curvature	$\nu_{Lf}, \nu_{Hf}, \nu_{Lm}, \nu_{Hm}$	Gender, elasticity of substitution	ρ
Non-pecuniary constraints	$\eta_{sL}, \eta_{sH}, \eta_{uL}, \eta_{uH}$	Group-gender productivity	$\theta_{Lf}, \theta_{Lm}, \theta_{Hf}, \theta_{Hm}$
		Group-gender elasticity of substitution	ρ_f, ρ_m
		Skill-group-gender productivity	$\theta_{sLf}, \theta_{sLm}, \theta_{sHf}, \theta_{sHm},$ $\theta_{uLf}, \theta_{uLm}, \theta_{uHf}, \theta_{uHm}$
		Skill-group-gender elasticity of substitution	$\rho_{Lf}, \rho_{Lm}, \rho_{Hf}, \rho_{Hm}$
		Land share	γ
<i>Castes with different endowments of land</i>			
		Caste group, land productivity	χ_L, χ_H
		Caste group, land elasticity of substitution	ρ_D
<i>Castes with different endowments of irrigated and unirrigated land</i>			
		Caste group, land productivity	χ_L, χ_H
		Caste group, land elasticity of substitution	ρ_D
		Caste group-irrigation, land productivity	$\chi_{r,L}, \chi_{r,H}, \chi_{n,L}, \chi_{n,H}$
		Caste group-irrigation, land elasticity of substitution	$\rho_{D,L}, \rho_{D,H}$


Note that augmenting the benchmark model changes the potential income but does not affect the first order conditions. To solve and estimate the model, however, it is no longer possible to split the parameters into two groups in Steps 1 and 2. This is because potential income depends on land income, for which we need to know all the parameters associated with the production function. Hence, we implement the estimation strategy as follows:

Step 1 Using observed education and FLFP, which gives us the labor input by caste, gender, and skill, as well as the total land endowment by caste group, predict wages in each population density interval for a particular choice of group 2 parameters. Perform a global search over all parameter values to find the set of parameters that minimizes the distance between observed and predicted wages.

Step 2 Given observed wages, group 2 parameter values identified in **Step 1** (which we need to compute the household's potential income), and a particular choice of group 1 parameters, solve equations (D.1), (D.2), (D.5), and (D.6)-(D.7) to derive $\tau_{sk}, \tau_{uk}, \xi_{kg}$ in each population density interval. Perform a global search over all parameter values to find the set of parameters that minimizes the distance between the data moments and the model moments with respect to FLFP and education.


Step 3 and Step 4 then proceed exactly as above.

Figure D1: Comparing education model and data moments, separately by gender and caste




Source: NSS, 1951 population census

Figure D2: Comparing wage model and data moments, separately by gender and caste




Source: NSS, 1951 population census

Figure D3: Female education: comparing the model and the data




Source: NSS, 1951 population census

Figure D4: Male education: comparing the model and the data




Source: NSS, 1951 population census

Figure D5: Female wages: comparing the model and the data




Source: NSS, 1951 population census

Figure D6: Male wages: comparing the model and the data




Source: NSS, 1951 population census

Figure D7: Model with land: comparison of parameters and counterfactuals



(a) Estimated parameters



(b) Counterfactuals