

NBER WORKING PAPER SERIES

GUARANTEED EMPLOYMENT IN RURAL INDIA:
INTRA-HOUSEHOLD LABOR AND RESOURCE ALLOCATION CONSEQUENCES

Jorge Luis García

Working Paper 33494
<http://www.nber.org/papers/w33494>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
February 2025

The views expressed herein are those of the author 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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2025 by Jorge Luis García. 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.

Guaranteed Employment in Rural India: Intra-Household Labor and Resource Allocation
Consequences

Jorge Luis García

NBER Working Paper No. 33494

February 2025

JEL No. I31, I32, J12, J13, O12, O15

ABSTRACT

I investigate the intra-household labor and resource allocation consequences of an employment guarantee targeting rural households in India. The guarantee insures household earnings, replacing women as added workers and shutting down a motive for saving. Despite sizable program-job take-up, the guarantee decreases participation in other working activities, and, thus, the labor force participation of married women and total time worked by their husbands. The guarantee accounts for up to 30% of a recent countrywide decrease in rural female labor force participation. Though it increases household consumption, the guarantee reduces the command of household earnings by women, and, thereby, their well-being.

Jorge Luis García

Department of Economics

Texas A&M University

2935 Research Parkway, Office 275

College Station, TX 77843

and NBER

jlgarcia@tamu.edu

1. Introduction

Women who participate in labor-market activities enjoy a degree of financial autonomy that participation in non-market activities does not provide them (see Kessler-Harris, 2003; Sen, 1990). Such autonomy determines their decision-making power and share of total resources within the household (e.g., Anderson and Eswaran, 2009; Blumberg and Coleman, 1989; Rahman and Rao, 2004). In India, the labor force participation of rural women decreased substantially during the last thirty years from an already low level,¹ suggesting a worsening of their economic conditions. I document that, by insuring household earnings, the Mahatma Gandhi Rural Employment Act shuts down a motive for precautionary savings and replaces rural married women as “added” or “insurance” workers. That is, the employment guarantee crowds out the participation of rural married women in the labor force, and, therefore, reduces their command of household earnings, intra-household share of consumption, and overall well-being.

In 2005, the Indian government enacted the Mahatma Gandhi Rural Employment Act. This employment guarantee insures the earnings of households whose individuals are willing to work in designated construction job sites in exchange for a daily minimum wage. It provides up to 100 job days per year per household, which can be freely distributed among adult members (Ministry of Rural Development, 2005a). The employment guarantee started between 2006 and 2008, depending on household location. Between 2012 and 2021, it provided at least one job day per year to an average of 81.5 million individuals (Government of India, 2022). Its primary statutory objective is to increase economic livelihood or security in rural areas. Its provision of jobs aims to insure households against the economic uncertainty typical of rural areas of India (Alik-Lagrange and Ravallion, 2018). Another of its objectives is empowering women by providing them with a job (Ministry of Rural Development, 2005b).

The act dictates that, in aggregate, women should hold at least one-third of employment-guarantee jobs at any time. This stipulation is non-binding. Between 2012 and 2021, 54% of employment-guarantee jobs were held by women. This aggregate statistic is inconclusive regarding the employment guarantee’s aim of empowering women by providing them with a job. The employment guarantee served an average of 44 million women with at least one job day per year during the referred period (Government of India, 2022). However, it likely shaped the labor-market decisions of millions more (India’s average rural population between 2012 and 2021 was almost 885 million; World Bank, 2022b). Indeed, by insuring

¹World Bank (2022a) reports a decrease in female labor force participation of 37%, from a participation rate of 30% in 1990 to a participation rate of 19% in 2021. As a result, India ranks 172 among the 181 countries for which this source reports female labor force participation in 2021.

household earnings, an employment-guarantee policy may compete with the role of women as added workers (i.e., it may compete with the role of women of supplying labor to provide insurance against economic shocks suffered by “primary” workers or “breadwinners”). Moreover, women working employment-guarantee jobs could have already been participating in the labor force before such jobs became available, implying that a large number of female employment-guarantee participants could have resulted from a shift across labor-market activities, rather than an increase in labor force participation.

I first analyze the impact of the employment guarantee on female labor force participation. For this analysis, I construct an individual-level, nationally representative sample using the repeated cross-sections of the Employment and Unemployment National Sample Survey (EU-NSS) that cover the period between 1999 and 2012. I combine these data with district-level variation in the timing of the employment guarantee and state-level variation in its intensity in an event-study framework. I find that the employment guarantee *reduces* the labor force participation rate of rural married women by four percentage points. This reduction is estimated across all observed labor-market activities; it is net of participation in employment-guarantee jobs.

Most of the decrease in female labor force participation observed in India during the last thirty years occurred during the period observed in the EU-NSS sample. Female labor force participation decreased from 35% to 27% between 2005 and 2012. This decrease was driven by rural married women, whose labor force participation decreased from 40% to 30% during this period. I do not argue that the employment guarantee drives the entirety of this decrease. My identification strategy recovers the average treatment on the treated, who are concentrated in seven of the 34 states and union territories analyzed. When population-weighting the negative impact on rural married women, I find that the employment guarantee accounts for up to 30% of the countrywide decrease observed for all rural women.

I use the Indian Human Development Survey (IHDS) to construct a longitudinal, nationally representative sample, including days worked by activity. In this sample, I estimate an impact on the labor force participation of rural married women essentially identical to that estimated in the EU-NSS. The estimate from the IHDS allows me to corroborate that the average treatment on the treated is driven by the average within-individual response to the policy. In both samples, I find that the employment guarantee does not affect the participation rate of rural married men.

For both rural married men and women, the employment guarantee shifts the distribution of days worked per year to the left. For women, the shift is such that the employment

guarantee not only increases the likelihood of not participating in the labor force but also the likelihood of working a relatively small number of days a year. For men, though there is no extensive-margin impact, there is also a shift to the left, which is an implication of the economic framework described below. The shift to the left in days worked occurs despite a sizable take-up of employment-guarantee jobs: on average, treated men and women work 4.2 and 10.3 days a year in employment-guarantee jobs. The negative impact on their annual workdays in other activities is larger in magnitude than the program take-up. The reduction in days worked across all activities, including employment-guarantee jobs, is 14.8 from a baseline average of 102 for married women and 9.5 from a baseline average of 217 for married men.

A basic economic framework explains the negative impact on labor force participation and days worked by rural married women and their husbands. If, together as a household, a woman and her husband plan their time spent working and are risk-averse towards the days of work available to them, they accumulate a buffer stock in anticipation of correlated negative shocks impeding their preferred (risk-free) work allocations. They finance this buffer stock by decreasing household consumption and days spent in non-market activities. Once the policy is in place, the household is permanently guaranteed a fixed number of annual workdays. This guarantee insures household earnings; it reduces the household's risk and thus its need to accumulate the buffer. While the households shocked in any given year take up employment-guarantee jobs, the average overall workdays across all activities decreases (i.e., all other households are not shocked and thus do not take up these jobs; on average, they reduce their overall days worked because they do not need to accumulate the buffer). The average number of days spent in non-market activities increases. So does average household consumption, as households prefer convex combinations of consumption and days spent in non-market activities rather than extremes.

Context-specific gender roles refine the implications regarding the decrease in participation of labor-market activities. In India, wives perceive their husbands as primary workers, and husbands prefer their wives not to work at all. Married men act like breadwinners or primary workers; their wives act as added, insurance, or secondary workers (Dean and Jayachandran, 2019; Eswaran et al., 2013; Jayachandran, 2021). Once the employment guarantee is in place, married men are likely to increase their non-market activities by decreasing their days worked (intensive margin); they are unlikely to quit the labor force (extensive margin). Married women see their role as added workers crowded out. For some of them, the reduced participation in labor-market activities includes completely quitting the labor force.

I test the economic framework’s implication that the employment guarantee increases household consumption. Using a cross-sectional dataset analogous to the EU-NSS and the longitudinal data of the IHDS, I apply the same empirical design as when assessing labor-market outcomes. I document that the employment guarantee *increases* monthly household consumption per capita by an average of 6% to 7% from a baseline average of 244 US dollars (2018, purchasing power parity). That is, the employment guarantee achieves its objective of providing rural households with economic security by insuring their earnings and thus increasing their consumption. By this standard, it reduces household-level absolute poverty. I also test the implication of a reduction in savings. While the available data on savings is limited to the IHDS and has caveats, I find an average reduction in the per-period flow of household savings of 99.6 from a baseline average of 237.5 US dollars (2018, purchasing power parity), as well as a small reduction in the ownership of physical assets.

By crowding out the labor force participation of rural married women, the employment guarantee reduces their contribution to household earnings. I argue that such reduction decreases their bargaining power and thus their intra-household share of resources. I assemble several imperfect but internally consistent pieces of evidence in favor of this argument. I first combine the quasi-experimental implementation of the employment guarantee with the estimation of a collective-household model (Chiappori, 1988, 1992) to determine how women and their husbands split the household-consumption gain generated by the employment guarantee. The relative gain for husbands is greater. The employment guarantee *reduces* the female intra-household share of resources, which has a one-to-one relationship with female bargaining power, by 9% from a baseline of 45% of the total household resources.

A decrease in female bargaining power limits domestic independence and increases intimate-partner violence (e.g., Anderson, 2021). I form a domestic-independence index in the longitudinal data from the IHDS. The employment guarantee decreases this index by an average of a third of a standard deviation, verifying the mechanism suggested by the structural estimates. Longitudinal data on the body-mass index (BMI), which has been used to measure the consequences of changes in within-household resource allocation (e.g., Calvi, 2020), further corroborate the structural estimates: I find that the employment guarantee has a substantial negative impact on BMI. Importantly, BMI is relevant as a measure of overall well-being, as it captures mental and physical health (Ackerson and Subramanian, 2008; Selvamani and Singh, 2018). The structural and reduced-form evidence indicates that, despite decreasing absolute household-level poverty, the employment guarantee makes women poorer within the household, hurting their overall well-being.

Related Literature. This paper relates to studies discussing the low level and recent

decrease in the labor force participation of women in India (e.g., Afridi et al., 2018, 2016; Bhargava, 2018; Desai and Joshi, 2019; Fletcher et al., 2017; Klasen, 2015). The reasons provided for the decrease include discrimination, education, increasing returns in home production, insufficient job creation, rising household earnings, rising male earnings, search frictions, and social norms. I find in the employment guarantee a plausibly causal reason that is new to the literature.

I argue that the gender roles of married individuals as either primary or secondary workers are fundamental in determining the impact of the employment guarantee on female labor force participation. This argument relates my findings to studies documenting that the cultural and institutional setting of a country determines its female labor force participation rate (Jayachandran, 2021). My explanation of why the employment guarantee reduces female labor force participation relates to studies assessing the non-market time allocation response of secondary workers to an improvement in household economic conditions (i.e., studies of the “added-worker effect,” e.g., Lundberg, 1985).

My structural results are consistent with studies documenting that a larger command of household earnings or assets by women increases their command of consumption decisions, household bargaining power, and intra-household share of resources (e.g., Attanasio and Lechene, 2014; Qian, 2008; Rangel, 2006). Calvi (2020) and Heath and Tan (2020) are related studies that focus on India. They argue that the possibility of inheriting property increases the intra-household bargaining power of women. Calvi (2020) finds that, as a consequence, their health improves (e.g., their BMI increases). Heath and Tan (2020) find that, as a consequence, they decide to participate more in the labor force.

The study by Field et al. (2021) helps to interpret my results and qualify the policy’s design. These authors experimentally alter the employment guarantee in the state of Madhya Pradesh. They allow a treatment group of women to receive payments from employment-guarantee jobs in their own private bank accounts. In the control group, the payments are directed to the male household head as is the *status quo* nationally. Field et al. (2021) find that the treatment of their experiment increases female labor force participation. I find that the national *status quo* implementation reduces female labor force participation in a setting where women are secondary workers, and where, even if they were to participate in employment-guarantee jobs, their payments would be directed to the male household heads. A joint interpretation of my findings and Field et al. (2021) indicates that the payment form is fundamental in achieving the policy’s aim of empowering women.

This paper also relates to studies evaluating India’s employment guarantee. After dis-

cussing my results, I provide an empirical comparison of the identification strategy in this paper to a common strategy in the literature (e.g., Azam, 2011; Imbert and Papp, 2015). This common strategy yields a positive (short-term) impact on rural wages. My strategy finds no long-term impact on rural wages. The difference is economically relevant because, in this paper, I argue that the impact of the employment guarantee on several outcomes is driven by its direct effect as insurance of household earnings. Other work argues that a primary channel is its increase of rural wages due to a general-equilibrium effect. Without a documented impact on human capital, the household insurance mechanism proposed in this paper is more plausible as an explanation for the labor force participation impacts, and accompanying increase in household consumption and decrease in household savings, than more indirect channels (e.g., general-equilibrium effects).

2. Data

Table 1 provides a self-contained summary of this section. I discuss the details next.

2.1 Labor-Market Analysis

Initial Samples. I construct two samples for analyzing labor-market outcomes. The first is a repeated cross-section. It is based on the seven cross-sections or rounds of the EU-NSS (Ministry of Statistics and Programme Implementation, 2020a) covering the period between 1999-2000 and 2011-2012. I pool the seven cross-sections to form a sample of women and men who were between 25 and 64 years old when they were surveyed. This sample includes all of the individuals who satisfy the age criterion independently of their household roles (i.e., head, child of head, or child-in-law of head). These individuals are observed across 34 of India’s 36 states or union territories.² In India, the administrative subdivision after a state is a district. The individuals in the sample are observed across 582 districts of India’s 640 districts in 2011-2012.³

The second sample is based on the two available rounds of the IHDS (Inter-University Consortium for Political and Social Research, 2011), which I use in longitudinal format. The

²I do not consider Delhi because it is mainly urban. For that reason, district classification of employment-guarantee availability is not available for this territory. Omitting Delhi has no consequences for replicating national patterns of labor force participation (see Section 4.1). I do not consider Ladakh because it is a very small in-conflict territory with a population of approximately 50,000 individuals (Office of the Registrar General and Census Commissioner, 2020).

³The districts observed vary across periods for two reasons. First, some districts are sampled in some rounds of the EU-NSS but not others. Second, the number of districts has increased over time from 593 in 2001 to 640 in 2011 (Office of the Registrar General and Census Commissioner, 2020). Despite their imbalance, the cross-sections are nationally representative. Appendix Table A.2 shows that the main results of the paper remain quantitatively similar and qualitatively unaltered when delimiting the sample to individuals from a balanced panel of districts encompassing the majority of the working sample described below.

Table 1. Summary of Analysis Samples

	(1)	(2)	(3)	(4)	(5)
	Labor-Market Samples		Consumption Samples		Female Well-being Sample
	Repeated Cross Sections (EU-NSS)	Panel (IHDS)	Repeated Cross Sections (HE-NSS)	Panel (IHDS)	Panel (IHDS)
Sampling					
Data Set of Origin	EU-NSS	IHDS	HE-NSS	IHDS	IHDS
Sampling Design	Seven cross-sections	Longitudinal	Seven cross-sections	Longitudinal	Longitudinal
Periods of Observation	1999-2000 to 2011-12	2004-2005 and 2011-2012	1999-2000 to 2011-12	2004-2005 and 2011-2012	2004-2005 and 2011-2012
Representativeness	National for each period	National for 2004-2005	National for each period	National for 2004-2005	National for 2004-2005
States in the Sample	34	21	34	21	21
Treatment	7	7	7	7	7
Control	27	14	27	14	14
Districts in the Sample	582	326	582	326	326
Observations					
Level	Individual	Individual (each observed twice)	Husband-wife pairs	Husband-wife pairs (each observed twice)	Married women (each observed twice)
Household Role	Head, child or child-in-law of head	Head, child or child-in-law of head	Husband is head	Husband is head, child or child-in-law of head	Female head, child or child-in-law of head
Age Profile	25 to 64 years old	25 to 64 years old (at both surveys)	Husband 25 to 64 years old	Husband 25 to 64 years old	25 to 64 years old (at both surveys)
Initial (Working) Sample Size or Number of Observations					
Urban	259,336 (218,717) men; 252,721 (214,451) women	8,105 men; 7,945 women	119,198 pairs	6,866 pairs	4,318 women
Rural	415,867 (376,896) men; 418,771 (380,775) women	18,514 men; 16,829 women	208,165 pairs	16,299 pairs	8,877 women
Outcomes					
	Labor force participation	Labor force participation, annual days worked by activity, daily wage	Household consumption per capita; private (non-shareable) consumption for each person in husband-wife pair	Household consumption per capita; household savings and livestock ownership	Domestic independence index, body mass index, height

Note: EU-NSS stands for Employment and Unemployment National Sample Survey. IHDS stands for Indian Human Development Survey. HE-NSS stands for Household Expenditure National Sample Survey. For the EU-NSS, the sample size of the working sample used for regressions appears in parenthesis. For all other samples, the initial and working samples are the same.

first round was in 2004-2005 and surveyed a nationally representative sample of households. The second round was in 2011-2012; it followed up with the households interviewed in the first round. I consider the individuals in the households observed in the two rounds. I construct a balanced panel using the same age and household role criteria that I used when forming the sample based on the EU-NSS. When reporting results using this sample, I display the number of individuals instead of the number of observations (individuals times periods). The individuals in the sample are observed across 326 districts of 21 states. Though more geographically limited than the EU-NSS, the IHDS is longitudinal at the individual level, while remaining nationally representative for the period 2004-2005. In both samples, I observe age, caste, religion, socioeconomic disadvantage, and marital status,⁴ and merge in district-level agricultural and state-level rain information from Government of India (2022).

Working Samples. I use the initial sample of the EU-NSS in annual (cross-sectional) format to describe the labor force participation in Section 4.1. In my regression analysis thereafter, I arrange the data quarterly to provide higher-frequency event studies. In this analysis, I discard observations from quarters with extremely thin cells to form the “working” sample. The annual participation patterns are virtually identical across the initial and working samples.⁵ In the IHDS, the regression analysis is annual. There is no difference between the initial and working samples.

Outcomes. In the EU-NSS sample, I construct labor force participation using a variable that indicates if an individual’s usual activity had been working or looking for work during the last year. The EU-NSS classifies individuals as having worked if they were self-employed or worked in a household enterprise, helped in a household enterprise (*paid* or *unpaid*), were salaried employees, were temporary or casual employees, or had any other type of employment. The EU-NSS classifies work in employment-guarantee jobs as temporary or casual employment (i.e., individuals participating in employment-guarantee jobs count as participating in the labor force).

In the IHDS sample, I construct the labor force participation variable as an indicator of having worked for money at least one day during the last year. This indicator is based on information on days worked by activity. I classify the working activities observed in four exhaustive and mutually exclusive categories: self-employment (*inside* or *outside* the

⁴I classify an individual as disadvantaged if they are Adivasi or Dalit (referred to as “scheduled castes and scheduled tribes”) or belong to “other backward classes (OBC).” All other individuals are non-disadvantaged in my classification (i.e., Christian, Jain, Jewish, Muslim, Sikh, or belonging to “forward castes”).

⁵For example, in the initial sample, the female labor force participation rate in the 1999-2000 cross-section is 39.782%, while it is 29.945% in the 2011-12 cross-section. The analogous rates for the working sample are 39.782% and 29.959%.

household), holding an agricultural job, holding a non-agricultural job, and holding a job provided by the employment guarantee (available in the second round of the survey, once the policy is in place). When analyzing days worked by activity as an outcome, I do not condition on participation (individuals not participating in an activity are assigned 0 days worked). In the IHDS sample, I also observe daily wages (earnings per day worked). I convert daily wages and all other monetary outcomes to 2018 purchasing-power-parity (PPP) dollars.

Missing Information. The number of observations in the most basic specifications that I present below corresponds to the number of individuals sampled. There are essentially no missing values in demographic and outcome variables. In some specifications, there are missing values due to missing district-level agricultural information or empty cells in specifications that interact age fixed effects of married women and their husbands. Missing values are, therefore, a minor issue not addressed empirically.

Summary Statistics. Panel a. of Table 2 summarizes baseline demographics in the EU-NSS and IHDS labor-market (working) samples and demonstrates alignment. Panel b. is based on rural women, the focus of my analysis. It provides baseline demographics by state treatment status and sector for 2004-2005 (before the employment guarantee) and 2011-2012 (after). State treatment status is one of the sources of variation in the implementation of the employment guarantee. For now, treatment states are the seven states with high levels of program availability, while control states are all the other states (they have low levels). More details on this definition are next in Section 3. For this table, I use the two cross-sections of the EU-NSS that coincide in timing with the two waves of the IHDS. While there is little difference across treatment and control states in the age and marriage profiles of their individuals, treatment-state individuals are clearly at socioeconomic disadvantage.

The treatment-control differences in socioeconomic disadvantage observed in Table 2 underscore the importance of designing an empirical strategy that acknowledges the non-random assignment of treatment across states. I rely on versions of difference-in-difference estimators in such design. I note that only two difference-in-differences in Column (11) differ statistically from 0 when using a significance level of 10%.⁶ This pattern suggests that

⁶Age is an exception in the EU-NSS, but the magnitude is small. For autumn crops, the difference-in-difference is sizable in both samples, even though it is not precisely estimated in the IHDS. This may be an outcome of the employment guarantee. Column (11) shows that agricultural production was larger in control-state districts compared to treatment-state districts before the employment guarantee. In treatment states, monsoon rains are less voluminous, and autumn crops depend on them. Thus, there is a greater need for the employment guarantee as insurance against adverse weather conditions (Fetzer, 2020; Johnson, 2009). The large and negative difference-in-differences in the production of autumn crops are likely a crowding-out outcome (Bahal, 2019). Winter crops help verify this because they rely on irrigation, not monsoon rains. Thus, the employment guarantee should impact winter crops less, which Panel b. of Table 2 confirms.

difference-in-differences estimators effectively address discrepancies between the treatment and control states.

2.2 Consumption Analysis

Samples. I construct two samples for analyzing consumption outcomes. The first is based on the HE-NSS (Ministry of Statistics and Programme Implementation, 2020b), which is identical to the EU-NSS in sampling characteristics. I pool the seven rounds of the HE-NSS that coincide in timing with the seven rounds of the EU-NSS described in Section 2.1.⁷ I construct a sample in which the observation level is the husband-wife pair of the male household head. In this construction, examples of demographic characteristics include age of the male household head and his wife. This construction differs from the construction of the sample based on the EU-NSS, which includes individuals independent of their household role. Despite this sample-construction difference, Panel a. of Appendix Table A.1 indicates that the basic demographics of the sample based on the HE-NSS closely align with those described in Table 2 for the sample based on the EU-NSS.

The second sample is based on the IHDS. I construct a longitudinal sample using all the observed husband-wife pairs, independently of the role of the husband in the household. I consider pairs that were observed in both rounds and construct a balanced sample using the same age criteria for the husbands that I use when forming the sample based on the HE-NSS. The IHDS consumption sample contains all of the married men in the IHDS labor-market sample. Given the very high husband-wife age correlation of 0.94, it also contains most of the married women. By construction, the IHDS consumption sample includes more husband-wife pairs per household than the HE-NSS consumption sample. This broader inclusion of husband-wife pairs allows me to verify that the restriction of only considering one husband-wife pair per household in the HE-NSS consumption sample does not introduce biases. For brevity, I refer to the husband-wife pairs in the HE-NSS and IHDS as households, despite the latter pairs originating from one of the potentially multiple pairs in a household. In both of the consumption samples, missing values are a minor issue not addressed empirically.

Outcomes. In both the HE-NSS and IHDS consumption samples, I observe total consumption of goods of the households where the relevant husband-wife pairs live. I construct the respective household consumption per capita. The HE-NSS also allows me to construct a composite of private (non-shareable) assignable consumption for women and their husbands.

⁷In the HE-NSS, I directly discard observations from quarters with extremely thin cells, and do not make a distinction between the initial and working samples in any of the analysis.

Table 2. Summary Statistics: Labor-Market Samples

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
	Panel a. Full (Pooled) Sample				Panel b. Sample of Rural Women							
	Men		Women		2004-2005			2011-2012			Diff-in-Diff	
	<i>Urban</i>	<i>Rural</i>	<i>Urban</i>	<i>Rural</i>	<i>Control</i>	<i>Treat</i>	<i>Diff</i>	<i>Control</i>	<i>Treat</i>	<i>Diff</i>	<i>(10)-(7)</i>	<i>p-value</i>
Individual-level												
<i>Age</i>												
EU-NSS	39.88	40.33	39.74	39.99	39.79	40.02	0.23	40.18	40.61	0.43	0.205	0.019
IHDS	41.90	41.95	41.80	41.69	38.03	38.59	0.56	45.08	45.48	0.41	-0.152	0.998
<i>Disadvantaged</i>												
EU-NSS	0.55	0.73	0.55	0.73	0.70	0.83	0.13	0.72	0.85	0.13	-0.004	0.982
IHDS	0.51	0.69	0.52	0.71	0.67	0.81	0.14	0.67	0.80	0.14	-0.001	0.998
<i>Literate</i>												
EU-NSS	0.87	0.66	0.70	0.38	0.37	0.28	-0.10	0.47	0.41	-0.06	0.038	0.996
IHDS	0.87	0.70	0.69	0.37	0.40	0.29	-0.11	0.41	0.30	-0.11	0.003	0.995
<i>Married</i>												
EU-NSS	0.87	0.90	0.86	0.88	0.88	0.86	-0.02	0.89	0.88	-0.02	-0.001	0.977
IHDS	0.89	0.92	0.83	0.84	0.87	0.88	0.00	0.80	0.81	0.00	0.000	0.998
District-level												
<i>Autumn Crops</i>												
EU-NSS		562.42		555.02	473.96	369.77	-104.19	966.14	448.08	-518.06	-413.867	0.005
IHDS		714.93		690.33	717.03	263.11	-453.92	974.19	371.38	-602.81	-148.891	0.992
<i>Winter Crops</i>												
EU-NSS		333.74		327.92	322.83	205.92	-116.92	473.06	327.00	-146.06	-29.139	0.989
IHDS		349.25		320.91	297.83	182.02	-115.80	423.95	265.65	-158.30	-42.505	0.999
State-level												
<i>Rain</i>												
EU-NSS	1.27	1.30	1.28	1.30	1.40	0.91	-0.49	1.40	1.03	-0.37	0.119	0.980
IHDS	1.38	1.34	1.40	1.35	1.47	1.17	-0.30	1.43	1.07	-0.36	-0.063	0.996
Observations												
EU-NSS	218,717	376,896	214,451	380,775	59,472	20,668		46,522	15,451			
IHDS	16,210	37,028	15,890	33,658	11,434	5,395		11,434	5,395			

Note: Panel a. displays the average or number of observations for the labor-market (working) samples described in Table 1. Panel b. is analogous in format to Panel a. except that it limits the sample to rural women for the years in the label and corresponding group of states (Control/Treat). Column (12) displays the state-clustered jackknifed wild-bootstrapped p -value associated with the null hypothesis of 0. **Variable definitions:** Disadvantaged: Adivasi and Dalit (“scheduled castes and scheduled tribes”) or “other backward classes.” All other individuals are non-disadvantaged in my classification. Married: currently married, as opposed to single (never married), divorced, or widowed. Autumn-crop production: Kharif crop yields in thousands of bushels per acre. Winter-crop production: Rabi crop yields in thousands of bushels per acre. Monsoon rain: monsoon rain in liters per square meter. Control/Treat: belonging to either the control or treatment states defined in Section 3.

All consumption variables are monthly.⁸ In the IHDS, I also observe a measure of per-month household (accumulation of) savings and an indicator of whether the household owns livestock, which I use as a proxy for physical assets.⁹

2.3 Female Well-Being Analysis

Sample. The IHDS collected well-being measures for a subsample of the married women in the sample described in Section 2.1. I construct a balanced panel based on this subsample.

Outcomes. I observe the binary responses to two sets of questions longitudinally. The first set of questions allows me to construct indicators for not agreeing with a woman in the community being beaten if she leaves the house without her husband’s permission, has an extramarital affair, brings no dowry to the marriage, neglects household chores, or is bad at cooking. The second set of questions allows me to construct indicators for women not needing permission from their husbands to go to the health center alone, visit a friend, or go to the store. I construct a “domestic-independence index” averaging the responses to these questions. I standardize this index to an in-sample mean of 0 and a standard deviation of 1. For this index, there is a sizable amount of non-response, which, as discussed below, qualifies the results based on it. In this sample, I also observe BMI and height.

3. The Mahatma Gandhi National Rural Employment Guarantee Act

The employment guarantee provides casual work in construction job sites. The work is casual because individuals who take up a job one day do not need to commit to additional work days. Payment is daily, at the minimum wage. Individuals perform low-skill tasks (e.g., moving piles of dirt). The employment guarantee provides up to 100 job days per household. Households are free to decide how these days are split between adult individuals. The primary objective of the employment guarantee is to “enhance the livelihood security in rural areas (Ministry of Rural Development, 2005a).” Ministry of Rural Development (2005b) states secondary objectives, which include generating infrastructure, protecting the environment, reducing rural-urban migration, and fostering social equity. Another of its secondary objectives is empowering women by promoting their labor force participation.

⁸The assignable, private good for women is a composite of the following items: sari (traditional female garment), hair oil, hair shampoo, hair cream, and sanitary pads. For their husbands, the corresponding composite good includes dhoti (traditional male trousers), lungi (traditional male sarong), shaving blades, shaving stick, razor, shaving cream, aftershave lotion, tobacco (and similar), paan, and alcoholic drinks. The construction of the composites uses all the goods that can be classified as assignable and private. In the 2004-2005 IHDS nationally representative sample, 2% of rural married women smoked tobacco or consumed similar intoxicants and 1% drank alcohol. For their husbands, the respective percentages are 39% and 21%.

⁹Household savings are constructed as income less consumption, all in monthly totals.

District-Level Implementation Phases. The large scale of the employment guarantee required a gradual implementation (Ministry of Rural Development, 2005a). The federal government mandated that certain districts had priority, determined by the presence of a Maoist insurgency, agricultural conflicts, and low human capital (Ministry of Rural Development, 2007). Other districts also had priority because they were classified as disadvantaged by an index constructed to advise social policy (Planning Commission, 2003). Prioritized districts were at a relative socioeconomic disadvantage by design. The employment guarantee came into existence earlier in prioritized districts: April 2006 for Phase-1 districts and April 2007 for Phase-2 districts. In all other districts (Phase-3 districts), it came into existence on April 2008.¹⁰

Treatment and Control States. Availability of employment-guarantee job sites and, therefore, aggregate (federal) employment-guarantee job provision is driven by seven states. Imbert and Papp (2015) and Klonner and Oldiges (2022) explain that supply-side program-availability factors (e.g., administrative capacity and experience in providing social programs) determine the difference between these seven states and the rest of India.

I label the seven states with greater program availability as “treatment states” and the remaining states as “control states.”¹¹ While program availability is greater in the treatment states, there was some provision in control states. I classify individuals according to the date on which the employment guarantee began in their district (treatment timing) and the treatment status of their state (treatment intensity). Importantly, all states had districts in either of the three implementation phases. Thus, I can make average comparisons between groups of districts belonging to either treatment or control states, within either of the implementation phases. Household-level information on employment-guarantee availability, take-up, and payment observed in rounds 66 (2009-2010) and 68 (2011-2012) of the EU-NSS labor-market sample allows me to clarify the meaning of state-level treatment status.

A household interested in participating in the employment guarantee usually registers to obtain a “job card,” which tracks annual workdays and information such as preferred payment method. Households aiming to participate or households that have participated

¹⁰The classification of districts by phase is available in Ministry of Rural Development (2010). The EU-NSS labor-market and HE-NSS consumption samples described in Section 2.1 include observations from 582 districts (186 belong to Phase 1, 121 belong to Phase 2, and 275 belong to Phase 3).

¹¹In the 2011 census, 25% of the Indian population inhabited treatment states and 75% control states (Office of the Registrar General and Census Commissioner, 2020). The treatment states are Andhra Pradesh, Chhattisgarh, Himachal Pradesh, Madhya Pradesh, Rajasthan, Tamil Nadu, and Uttarakhand. In samples based on the IHDS, the fourteen control states observed are a subset of the 27 control states observed in the EU-NSS and the HE-NSS. Section 2 documents that, despite this difference, the samples are comparable in average demographics, both when pooling states and when computing averages by treatment status.

in the employment guarantee have a job card. Registration is provided within two weeks, though the process is usually quicker (Ministry of Rural Development, 2005b). Generally, job-card rates speak to availability, while having a job card and seeking and obtaining an employment-guarantee job speaks to provision and take up.

Panel (a) of Figure 1 shows a substantial difference between the fraction of rural households having a job card in treatment and control states. This figure solidifies the greater availability in treatment states relative to control states: more than 60% of households in treatment states have come into contact with the program. Given that between 70% and 80% of households are disadvantaged and could benefit from the program, the suggested availability for potential beneficiaries is very large in treatment states. Uncertainty regarding program availability from a participant’s perspective seems minor in these states.


Figure 1 also describes employment-guarantee status among households with a job card (Panel b) and without a job card (Panel c). In treatment states, about 70% of households with a job card have ever sought and obtained an employment-guarantee job. The difference with control states is about 10 percentage points. In treatment states, the fraction of households that have ever sought and not gotten an employment-guarantee job is less than 10% (i.e., small in absolute terms and relative to control states). Panel (b) indicates greater availability and provision in treatment states. It shows that uncertainty about program provision is lower in treatment states. However, the non-provision rate of 10% for job seekers highlights that some uncertainty persists even after program implementation.¹²

Finally, Panel (d) displays the payment method for employment-guarantee jobs among households that have ever gotten one such job. Lack of payment seems to be a minor issue in control states and even a more minor issue in treatment states. Figure 1 indicates that the main treatment-intensity difference between treatment and control states is availability, rather than provision (i.e., job assignment, given a job is requested) or payment (i.e., job payment, given a job is performed). It indicates that, by 2009-2010, the employment guarantee was a reasonably reliable source of employment in treatment states. It also indicates that it was not as reliable of a source in control states, mainly because of lower availability.

Employment-Guarantee Relevance Among Working Activities. Panel (a) of Appendix Figure A.1 uses the IHDS rural labor-market subsample, in which I perform my analysis by working activity. It shows that the percentage of employment-guarantee annual

¹²Panel (c) is less informative because, in both groups of states, most households have never sought an employment-guarantee job when they do not have a job card. While about 20% of households without a job card in either treatment or control states have sought and not obtained an employment-guarantee job, this could be a result of the job card being processed at the time of the survey.

Figure 1. Employment-Guarantee Availability, Provision, and Take-Up by District-Level Implementation Phase and State-Level Treatment Status



Note: Panel (a) displays the fraction of rural households possessing an employment-guarantee job card. The fraction is displayed by district-level implementation phase and state-level treatment status. Panel (b) displays the fraction of households among those possessing an employment-guarantee job card that sought and obtained an employment-guarantee job during the last year, as well as the fractions of those who sought and did not obtain such employment or did not seek such employment. Panel (c) is analogous in format to Panel (b) for households not possessing an employment-guarantee job card. Panel (d) displays the fraction of households receiving payment for their employment-guarantee work by direct deposit (post-office or bank account, smart card), another form of payment (in person at gram sabha meeting, in the field, by a self-help group member, or any other), or have not been paid. The calculations for Panel (d) include households with and without employment-guarantee job cards. **Sample:** EU-NSS labor-market working sample delimited to 2009-10 and 2011-12 for Panel (a) and to 2009-10 for Panels (b) to (d).

workdays out of the annual workdays across all activities is large in treatment states, relative to the control states. In treatment states, women spend an average of 12.3% of their annual workdays in employment guarantee jobs in 2011-12, while men spend an average of 3.9%. In the control states, the analogous percentages are 2.2% and 1.7%.

Employment-Guarantee Wages. By law, the employment guarantee pays the minimum wage per day worked. However, the observed employment-guarantee wages may vary because minimum wages differ across and within states as dictated by local regulations (Chief Labour Commissioner, 2022). Geographic location (even within a state) and a worker’s skill determine the minimum wage.¹³

Panel (b) of Appendix Figure A.1 uses the 2011-2012 observations of the IHDS rural labor-market subsample.¹⁴ I compare the wages in the employment guarantee and in any other activity to the overall wage distribution. The wages in any other activity uniformly fit into the quintiles of the overall distribution by design. Employment-guarantee wages mostly fall into the second quintile of the overall wage distribution. My findings below indicate that this does not mean that wages available to individuals are higher in the employment guarantee than in their other jobs. Instead, geographic wage variation is such that the relatively high-wage employment-guarantee jobs are not available for those whose jobs pay wages at the bottom of the distribution.

Corroborating the Employment-Guarantee Information in the IHDS. Four exercises corroborate that the IHDS labor-market rural subsample replicates aggregate moments from official records of the employment guarantee available in Government of India (2022). First, 8.6% of all individuals in the sample participate in the employment guarantee, which is very close to the 9.2% participation rate in Government of India (2022). Second, among participants of employment-guarantee jobs, 52.3% are women in the sample while 51.3% are women in Government of India (2022). Third, 40% of participants belong to scheduled castes or tribes in the sample while 44.9% in Government of India (2022) belong to these groups. Fourth, the average days in employment-guarantee jobs among those participating in such jobs for at least one day in treatment states is 37 in the sample and 36 in Government of India (2022). In control states, the averages are 33 and 31.¹⁵

¹³I do not exploit this variation in the analysis below because the fixed effects and controls in the main specifications absorb it.

¹⁴I use the male subsample to avoid the standard sample selection issue in the observation of female wages, which may be substantial in India due to its low female labor force participation rate. For men, sample selection is a minor issue (see Section 4.1).

¹⁵The first three comparisons use aggregate national statistics reported in Government of India (2022) for 2012-2013. The fourth comparison uses aggregate statistics by state reported in the same source for 2018-2019. The earliest period for which this source reports aggregate national statistics is 2012-2013, while the

Implementation Issues. Dutta et al. (2012, 2014) document that the job days demanded by individuals were larger than the job days supplied by the government during the initial years of the employment guarantee. The authors argue that the gap was due to provision-capacity constraints and that it was larger in poorer states. Imbert and Papp (2014) and Banerjee et al. (2020) argue that the gap diminished over time. The evidence in Dutta et al. (2012, 2014) indicates that the provision was below the maximum of 100 days across states in the initial years of the employment guarantee; it also indicates that implementation improved over time. Even if the maximum remains below 100 annual days due to capacity constraints, the argument throughout the paper does not change. I assume that, when making employment choices, individuals take the maximum days available as exogenously determined by the authorities implementing the program (just as they would take 100 days).

Funding and Corruption. I analyze the employment guarantee as implemented nationally. The (federal) Department of Rural Development funds 75% of the employment-guarantee operation costs and all of the wages of its beneficiaries (Ministry of Rural Development, 2005a). When flowing from the federal to the local level, the funds need to pass through various bureaucratic layers. Banerjee et al. (2020) document pervasive funding leakage during the first years of the employment guarantee. They implement a field experiment in the state of Bihar and find that a transparency reform reduces leakage. A similar reform to that in Bihar took place nationally in 2011.¹⁶ No evidence indicates that corruption compromises the individual or household average treatment effects discussed below. However, funding leakage necessarily implies that, in practice, the program is implemented with less intensity than originally planned. In that case, the estimates below are absolute-value lower bounds of an uncorrupted implementation.

4. Labor Force Participation and the Employment Guarantee

4.1 Labor Force Participation Before and After the Employment Guarantee

I first use the EU-NSS labor-market initial sample to describe the aggregate context of labor force participation before and after the employment guarantee. For each of the cross-sections composing this sample, I plot the labor force participation by sex and sector in Panels (a) and (b) of Figure 2. For women, the participation rate remained virtually constant between 1999-2000 and 2005-2006. After that, it decreased by eight percentage points.¹⁷ This substantial

earliest period for which it reports aggregate statistics by state is 2018-2019. Trends in national aggregate statistics reported in this source indicate stability between 2012-2013 and 2018-2019.

¹⁶The reform mandated live updates on Government of India (2022) about funding and participation.

¹⁷World Bank (2022a) reports that the additional decrease after 2012 is minor relative to the decrease observed in Figure 2. I do not use the three EU-NSS rounds before 1999-2000 because they do not contain geographic identifiers, an essential component of my analysis. In Appendix Figure A.2, I expand the sample

decrease occurred while the female labor force participation rate was increasing worldwide, and the male labor force participation rate in India barely changed, as displayed in Panel (b) of Figure 2. The decrease for women occurred from an already low rate. In 2005, India ranked 158 among the 181 countries for which the World Bank (2022a) documents female labor force participation. In 2012, it ranked 168. Similarly, it ranked 165 in the female-to-male labor force participation rate in 2005 and 172 in 2012.

Panel (a) of Figure 2 shows that the overall decrease observed for women is driven by those who are rural. Panel (c) of Figure 2 shows that, among them, the decrease is much more pronounced for the married. Rural married women are the main subsample in the rest of the paper.¹⁸ I group them into their district-level implementation phase and state-level treatment status to plot their quarterly participation rates in Panel (a) of Figure 3, shifting the time axis from calendar-year quarters to event time or quarters after the implementation of the employment guarantee.¹⁹ Panel (b) plots the corresponding treatment-control difference for each phase. That is, it displays the rate for those who live in a Phase-1 district that belongs to a treatment state less the rate for those who live in a Phase-1 district that belongs to a control state. It also displays the analogous rate differences in Phases 2 and 3. For Phases 1 and 2, this preliminary, raw analysis indicates that the employment guarantee could be a cause for the decrease observed after 2005.²⁰ A formal analysis is next.

4.2 Frameworks for Micro-Data Analysis

Event-Study Framework for Repeated Cross-Sectional Data. Let y_{ig} indicate labor force participation for woman i during event quarter g . The index i suffices to describe her district-level implementation phase and her state-level treatment status.

I use the index $p \in \{1, 2, 3\}$ to label Phase- p coefficient estimands. I partition the

to include the three rounds observed before 1999-2000 and show that the male and female labor force participation rates barely change between 1983 and 1999-2000.


¹⁸Except for a brief analysis of marital status as an outcome or comparisons to urban married women, the remaining analysis is based on rural married women, who are the vast majority of women in rural India. Analysis of unmarried rural women is unreliable due to a small number of observations in either of the samples considered. Comprehensive investigation of this minority is outside the scope of this paper.

¹⁹Appendix Figure A.4 is analogous in format to Panel (a) of Figure 3 without the shift to event-time quarters. I define April 30, 2006, as the start of the employment guarantee in Phase-1 districts; April 30, 2007 in Phase-2 districts; and April 30, 2008 in Phase-3 districts. I use the 30th as the starting day to make the conversion between calendar-year and event-year quarters straightforward and visually clear. The precise starting day in April is of minor importance (household interviews happened throughout the year; only a handful of them occurred during a given day of April).


²⁰Appendix Figure A.8 is analogous to Figure 3 for urban married women. There are no post-implementation trend breaks for either phase. I am able to calculate event time for urban women because most districts have urban and rural areas.

Figure 2. (Raw) Labor Force Participation in India


(a) Women by Sector



(b) Men by Sector



(c) Rural Women




(d) Rural Men



Note: Panel (a) displays the fraction of women who participated in the labor force during the twelve months prior to the interview conducted in the survey round corresponding to the year in the horizontal axis. The calculation includes married and unmarried (never married, separated, divorced, or widowed) women who were between 25 and 64 years old during the survey. Panel (c) breaks out the labor force participation rate of the rural women in Panel (a) into the participation rates of those married and unmarried. It also breaks out the labor force participation rate of rural married women into the participation rates of those disadvantaged and non-disadvantaged, as defined in Table 2. Panels (b) and (d) are analogous in format to Panels (a) and (c) for men. **Sample:** EU-NSS labor-market initial sample.

Figure 3. (Raw) Labor Force Participation of Rural Married Women by District Phase and State Treatment Status



Note: Panel (a) displays the labor force participation rate of rural married women by district-level implementation phase and state-level treatment status. The rates are displayed for each observed event-time quarter or quarter after the implementation of the employment guarantee. Panel (b) displays the quarterly treatment-control differences in Panel (a) by phase. In Panel (b), I subtract the treatment-control difference in the closest period to 0 before implementation (reference period) from each of the quarterly treatment-control differences. The treatment-control difference in the reference period is thus set to 0 and appears in the plot without a confidence interval. Panel (c) displays the quarterly weighted average across phases of the treatment-control differences in Panel (c), where the weights are the fractions of the population corresponding to each phase during each quarter. Panels (b) and (c) display the 95% confidence interval based on the jackknifed wild-bootstrapped distribution clustered at the state level for each treatment-control difference. Panels (b) and (c) display the average weighted treatment-control difference before and after implementation, relative to the reference period. The jackknifed wild bootstrapped p -value clustered at the state level associated with the null hypothesis of 0 accompanies each of these differences. **Sample:** Rural married female subsample of the EU-NSS labor-market working sample.

individuals in the EU-NSS labor-market working sample by district-level implementation phase and estimate the following model in each subsample:

$$y_{ig} = \sum_{j \in \mathcal{J}^p} \tau_j^p \cdot \mathbf{1}[g = j]_g + \sum_{j \in \mathcal{J}^p} \gamma_j^p \cdot \mathbf{1}[i \text{ lives in a treatment state}]_i \cdot \mathbf{1}[g = j]_g + \varepsilon_{ig}, \quad (1)$$

where $\mathbf{1}[\cdot] = 1$ if the statement in brackets is true and $\mathbf{1}[\cdot] = 0$ otherwise and ε_{ig} is an error term. τ_g^p is the average of y_{ig} in quarter g for districts in control states implementing the employment guarantee in Phase p , and γ_g^p is the corresponding treatment-control difference. Panel (a) of Figure 3 displays the estimates of τ_g^p and the corresponding treatment-state averages (i.e., estimates of $\tau_g^p + \gamma_g^p$). The estimates are displayed for the observed quarters, indexed by $\mathcal{J}^p$, which differ by phase subsample. Panel (b) displays estimates γ_g^p .²¹

Difference-in-Difference Estimands. Pre-policy difference between districts and states, evident in Table 2 and Figure 3, make specifications accounting for district fixed effects necessary explorations. In such case, I modify Equation (1) to

$$y_{ig} = \nu_d + \sum_{\substack{j \in \mathcal{J}^p \\ j \neq \underline{j}^p}} \tau_j^p \cdot \mathbf{1}[g = j]_g + \sum_{\substack{j \in \mathcal{J}^p \\ j \neq \underline{j}^p}} \gamma_j^p \cdot \mathbf{1}[i \text{ lives in a treatment state}]_i \cdot \mathbf{1}[g = j]_g + \varepsilon_{ig}, \quad (2)$$

where $\underline{j}^p \in \mathcal{J}^p$ with $\underline{j}^p \leq 0$ is the quarter closest to 0 observed for women who live in districts implementing the employment guarantee in Phase p and ν_d represents a district fixed effect (which subsumes the corresponding state fixed effect). In this case, γ_g^p is the quarterly (conditional) labor force participation rate for individuals who reside in Phase- p districts located in treatment states less the same rate for those who reside in Phase- p districts located in control states. I exclude a reference period ($\underline{j}^p$) because the fully saturated version of Equation (2) is not identified. This specification implies that the coefficients γ_g^p are relative to the treatment-control difference in the reference period (i.e., they are difference in differences).

Parameter of Interest and Identification Assumptions. My empirical design is akin to that in Sun and Abraham (2021). The coefficients γ_g^p are estimands of the quarterly average treatment on the treated (ATT). Their estimation is based on phase-wise treatment-control comparisons, or, more precisely, a “high-dose to low-dose of treatment” comparison. These comparisons are an absolute-value lower bound of the ideal “high-dose to no-dose of treatment” comparison (Heckman et al., 2000), which is a caveat. I can only identify and

²¹I follow usual practice and subtract from each estimate of γ_g^p the estimate of γ_g^p for which g is the largest and $g \leq 0$. This normalization to a reference period makes the interpretation of the coefficient estimates from Equations (1) and (2) identical (i.e., they are difference in differences).

estimate the referred lower bounds.²²

The coefficient γ_g^p identifies the ATT under two assumptions: no anticipation and parallel trends. The two assumptions together imply the testable implication of no expected difference in levels before implementation: $\gamma_g^p = 0$ for $g < \underline{j}^p$. If this implication holds, parallel trends before implementation hold, which favors parallel trends after implementation. Panel (b) of Figure 3 displays evidence in favor of the implication. It is based on estimates of γ_g^p from Equation (1). Below, evidence from the estimation of Equation (2) also favors this implication.

Aggregating the ATT Across Phases. To obtain a quarterly national estimate of the ATT, I aggregate γ_g^p across phases using the estimator

$$\text{WDiD}_g := \sum_{p=1}^3 [\text{fraction of individuals in Phase-}p \text{ districts}]_g \cdot \gamma_g^p, \quad (3)$$

which is a quarterly population-weighted sum of the phase-wise difference-in-difference ATT estimators.²³ In practice, estimating the WDiD_g requires estimates of γ^p and the quarterly population weights, which I estimate using their sample counterparts. Panel (c) of Figure 3 displays estimates of WDiD_g corresponding to estimates of γ_g^p from Equation (1).

Aggregating the ATT Over Time and Phases. To obtain a static summary of the ATT, I aggregate the ATT over time and phases. Minor adaptations of Equations (1) and (2) allow me to estimate the ATT before and after the employment guarantee, relative to the reference period $\underline{j}^p$, simply by lumping all the quarters before $\underline{j}^p$ as one period and all the quarters after $\underline{j}^p$ as another period. This allows me to estimate the ATT for each phase relative to $\underline{j}^p$ for either lumped period. Let γ_{after}^p denote the treatment-control difference after $\underline{j}^p$. The aggregate ATT is

$$\text{WDiD}_{\text{after}} := \sum_{p=1}^3 [\text{fraction of individuals in Phase-}p \text{ districts}]_{\text{after}} \cdot \gamma_{\text{after}}^p, \quad (4)$$

which is my main estimate in the EU-NSS labor-market working sample. The $\text{WDiD}_{\text{before}}$ is

²²The ATT in this paper is an “intent to treat” parameter. I aim to understand the impact of the sole existence of the insurance provided by the employment guarantee. Other parameters more directly focus on the actual take-up of employment-guarantee jobs. These parameters are not the focus of this study.

²³If only districts of Phase 1 are observed in a given quarter, then the point estimate is not weighted; it is simply the Phase-1 estimate. If only Phase-1 and Phase-2 districts are observed in a given quarter, the weights are relative to the total number of observations in these two phases. Similar definitions apply when districts of any other pair of phases are observed.

analogous in definition. Figure 3 displays estimates of these aggregate ATTs based on the adapted or lumped version of Equation (1).

Difference-in-Difference Framework for Longitudinal Data. In the IHDS labor-market sample, I observe individuals longitudinally: once in 2004-2005 (before the start of the employment guarantee) and once in 2011-2012 (after). These periods are either an entire year before or an entire year after the employment guarantee starts across implementation phases, ameliorating concerns related to variation in the timing of treatment and heterogeneity when considering an unweighted estimator that pools the observations from all of the districts (de Chaisemartin and d’Haultfoeuille, 2020). I thus estimate the following basic difference-in-difference model (two-period, two-treatment-status regimes) in this sample:

$$y_{ig} = \nu_i + \tau_{\text{after}} + \gamma_{\text{after}} \cdot \mathbf{1}[i \text{ lives in a treatment state}]_i \cdot \mathbf{1}[g = \text{after}]_g + \varepsilon_{ig} \quad (5)$$

pooling phases, where I reuse the notation defined above. $\gamma_{\text{after}} := \text{DiD}$ is a standard difference-in-difference estimator; it identifies the aggregate ATT under the assumption of parallel trends. I cannot provide the standard evidence in favor of this assumption in the IHDS sample. I rely on the event-study framework for justification. In this case, however, I am able to include an individual fixed, ν_i , which subsumes the corresponding district and state fixed effect. DiD is an appealing estimator for its simplicity and reliance on within-individual variation. It does not use the district-level variation in the timing of treatment. It is the (conditional) average treatment-control difference in the within-woman response to the employment guarantee between 2004-2005 and 2011-2012.

Inference. The inference throughout the paper is based on jackknifed wild bootstrap p -values as is recommended for settings with a “small” number of clusters and a “large” number of observations (Cameron et al., 2008; Canay et al., 2021). These p -values are associated with the null hypothesis of 0 (no impact). They account for sampling variation in all estimation stages—e.g., sampling variation in both the phase-wise average treatment on the treated and the weights in Equation (3). They are clustered at the state level, which is the widest level of treatment assignment, and, thus, the clustering recommended in (Abadie et al., 2022). The 95% confidence intervals displayed are based on the inversion of the corresponding p -values.

4.3 Main Impact: Rural Married Female Labor Force Participation

Figure 3 visually illustrates the identification argument and estimation of the main estimate of the ATT. In Panel (c), the absence of pre-policy treatment-control differences relative to the difference in the reference period is evident for the three phases, supporting the no-

anticipation and parallel-trends assumption. The post-implementation trend break is evident for Phases 1 and 2. It is not for Phase 3. The post-implementation aggregate estimate of the ATT across phases based on the $WDiD_{after}$ estimator in Equation (4) amounts to a decrease of 2.1 percentage points for rural married women (no-impact p -value = 0.001).


Figure 3 is based on the raw (no fixed effects or controls) specification in Equation (1). Conditioning on district fixed effects, which subsume state fixed effects, is a natural option given the employment-guarantee prioritization of disadvantaged districts, as well as the difference in employment-guarantee intensity and labor force participation by state-level treatment status. Equation (2) allows for that. Age and spouse-age fixed effects are also a natural conditioning set given the standard life-cycle profile of labor force participation. Agricultural and weather controls are another natural conditioning set, given that they are obvious determinants of aggregate rural labor force participation and given some of the treatment-control imbalances observed in Table 2. The additional fixed effects and controls are straightforward to incorporate into Equation (2).

Panel (a) of Figure 4 displays estimates of Equation (2) for each phase. It is based on the most comprehensive specification in the paper, which includes district fixed effects, age fixed effects, spouse-age fixed effects, and the district-level and state-level controls in Table 2 (entered into the equation linearly). Panel (b) displays the corresponding quarterly aggregate ATT estimates across phases. Evidence in favor of no anticipation and parallel trends is clear. The post-implementation aggregate estimate of the ATT amounts to -0.038 (no-impact p -value = 0.006).

Unlike the raw specification, the quarterly ATT differs from 0 with 95% confidence for all three phases by the end of the window of observation, when the program was most effectively implemented (see Section 3). A caveat that arises from comparing Panel (a) of Figure 3 to Panel (a) of Figures 4 is that the ATT from Phase 3 only differs from 0 statistically at standard significance levels conditionally and in the long term. Thus, results from this phase are more sensitive to specification than results from other phases. That said, upon the inclusion of district fixed effects, event-study estimates from alternative specifications (with or without age fixed effects, spouse-age fixed effects, or controls) are very similar to those in Figure 4 (see Appendix Figures A.5 to A.7).

Panel a. of Table 3 summarizes the estimation of these event studies using the post-implementation aggregate ATT, building from a specification with only district fixed effects (Column 1) to the specification displayed in Figure 4 (Column 4). Panel b. of Table 3 displays their counterpart estimates based on the DiD estimator of Equation (5), which

Figure 4. (Conditional) Labor Force Participation of Rural Married Women and the Employment Guarantee, Main Event Study



Note: Panel (a) displays estimates of γ_g^p for each quarter g and phase p based on Equation (2), including district, age, and spouse-age fixed effects, as well as the district-level and state-level controls in Table 2 (entered into the equation linearly). The estimates displayed are the (conditional) quarterly labor force participation rate in Phase- p districts located in treatment states minus the analogous rate for Phase- p districts located in control states. These treatment-control differences are relative to the difference in the closest quarter before implementation (reference period). The treatment-control difference in the reference period is thus set to 0 and appears in the plot without a confidence interval. Panel (b) displays the population-weighted average of the γ_g^p estimates in Panel (a) based on Equation (3). Both panels display the 95% confidence interval based on the jackknifed wild-bootstrapped distribution clustered at the state level for each treatment-control difference. Both panels display the average weighted treatment-control difference across phases before and after implementation, relative to the reference period, based on Equation (4). The jackknifed wild bootstrapped p -value clustered at the state level associated with the null hypothesis of 0 accompanies each of these differences. **Sample:** Rural married female subsample of the EU-NSS labor-market working sample.

relies on the IHDS labor-market working sample. I estimate specifications analogous to those considered when estimating Equation (2), except that individual fixed effects replace district fixed effects. Estimates are remarkably similar across the samples, reinforcing that within-individual policy responses drive them. I discuss the aggregate relevance of their magnitude in Section 4.4.

Rural Married Men. Columns (5) to (8) of Table 3 show that the employment guarantee has no impact on the labor force participation of rural married men. Section 5 explains this lack of impact in conjunction with the negative impact on rural married women.

Placebo Tests. Before discussing identification threats, I provide falsification tests in two placebo subsamples. The first subsample includes rural, married, and non-disadvantaged women. The second includes urban married women. The employment guarantee should have little to no impact on the women in the first subsample, as its jobs should be unattractive to them given the minimum-wage stipulation. A caveat in the construction of this subsample is that the definition of disadvantaged is coarse (i.e., it is only based on religion and caste). Some women classified as non-disadvantaged could thus participate in the employment guarantee. Yet, the disadvantaged should drive the impact among rural married women. The employment guarantee should also have no impact on the women in the second subsample because its jobs are not available in urban areas. By construction, the placebo subsamples and the subsample of rural married women differ in observed and unobserved characteristics. However, they allow for a minimal or basic falsification test. Appendix Table A.3 displays the impacts for both placebo subsamples, based on the EU-NSS and the IHDS labor-market samples. It indicates no impact in both cases.²⁴

Identification Caveats and Threats. The estimates based on the EU-NSS labor-market sample partly exploit within-district policy variation, which is more granular than the within-state policy variation on which the estimates based on the IHDS labor-market sample rely. In either sample, the parameter estimated is the ATT. In the former sample, I am able to document evidence in favor of the identification assumptions, while in the latter, I am able to document that the average within-woman policy response drives the impact on labor force participation. Ideally, such a response could be documented in a setting that exploits granular policy variation and is observed for multiple pre-implementation and post-implementation periods.

²⁴The results in Appendix Table A.3 imply that disadvantaged individuals, the *de facto* targets of the employment guarantee, drive the impact among rural married women (which contains both the disadvantaged and the non-disadvantaged). Quarterly event-study estimates are very imprecise, likely due to the small sample size, but the aggregate ATT estimates are precise at indicating a lack of impact.

Table 3. Labor Force Participation of Rural Married Women and Men and the Employment Guarantee, Estimates of the Average Treatment on the Treated

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
<i>Panel a. Data: EU-NSS; Year Span: 1999-2000 to 2011-2012; Estimator: WDiD</i>								
	Women				Men			
Fixed Effects	Dist	Dist, Age	Dist, Age	Dist, Age, Spouse Age	Dist	Dist, Age	Dist, Age	Dist, Age, Spouse Age
Controls	No	No	Yes	Yes	No	No	Yes	Yes
Estimate	-0.037	-0.037	-0.038	-0.038	-0.001	-0.001	0.000	0.000
(<i>p</i> -value)	(0.001)	(0.001)	(0.005)	(0.006)	(0.825)	(0.759)	(0.950)	(0.974)
Baseline Treatment Mean	0.644	0.644	0.643	0.643	0.976	0.976	0.976	0.976
Observations	333,957	333,957	301,881	301,790	333,105	333,105	300,631	300,589
<i>Panel b. Data: IHDS; Year Span: 2004-2005 and 2011-2012; Estimator: DiD</i>								
	Women				Men			
Fixed Effects	Indv	Indv, Age	Indv, Age	Indv, Age, Spouse Age	Indv	Indv, Age	Indv, Age	Indv, Age, Spouse Age
Controls	No	No	Yes	Yes	No	No	Yes	Yes
Estimate	-0.037	-0.034	-0.042	-0.039	-0.004	-0.005	-0.003	-0.004
(<i>p</i> -value)	(0.030)	(0.032)	(0.032)	(0.043)	(0.658)	(0.662)	(0.707)	(0.596)
Baseline Treatment Mean	0.677	0.677	0.675	0.675	0.963	0.963	0.962	0.962
Observations	13,375	13,375	13,154	13,142	16,328	16,328	16,078	16,050

Note: Column (1) of Panel a. displays details from the estimation of the aggregate average treatment on the treated based on Equation (4) (i.e., on the WDiD estimator) for rural married women. The required estimates of the average treatment on the treated for each phase are based on Equation (2). Columns (2) to (4) are analogous in format to Column (1). Their only difference is the inclusion of additional fixed effects or controls. The controls are the district-level and state-level controls in Table 2 (entered into the equation linearly). Columns (5) to (8) are analogous in format to Columns (1) to (4) for rural married men. Panel b. is analogous in format to Panel a. The estimate of the aggregate average treatment on the treated is based on Equation (5) (i.e., on the DiD estimator). Panel b. is based on longitudinal data rather than repeated cross-sections. It thus replaces district (Dist) with individual (Indv) fixed effects. For each estimate, the state-clustered jackknifed wild-bootstrapped *p*-value associated with the null hypothesis of 0 is displayed in parentheses. **Sample:** Rural married female (left) and male (right) subsamples of the EU-NSS (Panel a.) and IHDS (Panel b.) labor-market working samples.

Identifying the parameter of interest requires a time-invariant classification of the district and state where individuals reside. Rural-rural, across-district migration would compromise my identification strategy and could be a symptom of anticipative behavior. For example, individuals in Phase-3 districts could migrate to Phase-1 districts to obtain a job before the employment guarantee is available in their district of residence. This concern is not first-order because rural-rural, across-district migration is empirically irrelevant during the period that I analyze (Imbert and Papp, 2019). More generally, there is a low level of permanent migration in India (e.g., Munshi and Rosenzweig, 2009; Topalova, 2010).

I mainly focus on married women for three reasons: (i) they are the majority of women in rural India; (ii) they drive the aggregate decrease in female labor force participation; and (iii) the employment guarantee targets households and their adult members, which, in rural India, generally include married women and their husbands. This focus could be problematic. For example, the employment guarantee could improve men as marital prospects. Rural Indian women, whose probability of specializing in home production increases when getting married (Afridi et al., 2018), could perceive this improvement and increase their marriage rate. Marriage would then mediate the negative impact on the labor force participation of married women. Appendix Table A.4 is analogous in format to Table 3. The outcome is “being married,” as opposed to single (never married), divorced, or widowed. The corresponding event-study figures are in Appendix 3. They show no pre-trends or trend-breaks after the policy. The resulting estimates of the aggregate ATT are precisely estimated at 0 for both the male and female subsamples of the EU-NSS and IHDS labor-market working samples, discarding marriage as a relevant mediator.

The downward aggregate trend in female labor force participation is the most crucial identification threat. My design exploits district-level variation in the timing of implementation and state-level variation in intensity to tease out the differential decrease in treatment states from the aggregate trend. A visual walk through this identification argument is in Figure 3. The absence of pre-event trends supports the identification argument, though it does not make it definitive. Other reasons for the decrease in female labor force participation in the literature are mainly observational and descriptive, making it difficult to generate testable implications that could reinforce my identification argument. However, Chikermane (2018) provides a review of the seventy most important national policies in India during the period between 1947 and 2017. He does not discuss any policy that could be confounded with the employment guarantee because of its district-level timing or state-level intensity, or any policy targeting the labor force participation of rural women.

Finally, I note that estimates of the same parameter relying on different sources of

policy variation yield virtually identical estimates. They are based on the estimators that, in this empirical application, effectively difference out pre-policy differences between regions of treatment in the dependent variable of interest, as documented in Figures 3 and 4, and other observed characteristics, as documented in Table 2. The variation I exploit allows me to provide a plausible causal reason for the decrease in labor force participation. The studies surveyed before, when suggesting a reason for the decrease, are not based on sources of plausibly exogenous variation. An additional source of identification support is that, below, I pose an economic argument to explain the decrease in female labor force participation that generates testable implications regarding other household behavior (e.g., consumption and savings). I test and find support for these additional implications in Section 5.

4.4 Aggregate Relevance of the Impact on Female Labor Force Participation

I use the EU-NSS labor-market sample to illustrate the aggregate relevance of the main result (i.e., the negative impact on the labor force participation of rural married women). The calculation weights the aggregate ATT (over time and phases) by the proportion of women for whom the ATT applies, (i.e., rural married women in treated states) and divides the resulting weighted ATT by the aggregate decrease observed for all rural women in Figure 2. I perform this calculation for the aggregate decrease between 2005-06 and 2011-12 (total decrease) and the aggregate decrease between 2005-06 and 2007-08 (immediate decrease), for all the ATT estimates in Table 3. I find that the ATT is relevant in aggregate: a policy targeting about 22% of the population (rural married women in treated states) accounts for up to 30% of the immediate decrease and 10% of the total decrease.

5. Why Does a Household-Level Employment Guarantee Decrease the Labor Force Participation of Rural Married Women?

Section 3 documents substantial take-up of employment-guarantee jobs, particularly among rural married women. The negative impact on their labor force participation is, therefore, puzzling. To clarify this apparent contradiction, I estimate the impact of the employment guarantee on days worked across mutually exclusive categories of work using the DiD estimator in Equation (5) and longitudinal data on days worked by activity from the IHDS labor-market sample.²⁵ Panel (a) of Figure 5 displays the estimates, indicating an annual average treatment-control difference in the take-up of employment-guarantee jobs of 10.3

²⁵I estimate Equation (5) using the most complete specification, which includes individual, age, and spouse-age fixed effects, as well as controls. Days worked in employment-guarantee jobs are coded as 0 for all observations in 2004-2005 (before implementation). Appendix Figures A.11 to A.13 are analogous to Figure 5, using other specifications. Results remain aligned across the different specifications.

days. However, the employment guarantee also leads to average annual reductions in both agricultural and non-agricultural work of 12.2 and 13.4 days, from baseline averages of 40.8 and 20.2. These declines outweigh the employment-guarantee take-up, resulting in a net average annual reduction of 14.8 days from a baseline of 101.8.²⁶

The average annual decrease of 14.8 days combines both the extensive-margin and intensive-margin responses to the employment guarantee, as annual days worked are recorded as 0 for those who do not work. To clarify these responses, I estimate the impact of the employment guarantee on each of the indicators labeled in Panel (b) of Figure 5. The first estimate shows an increase in the probability of working 0 days a year of 3.9 percentage points, which corresponds to the negative impact on participation reported in Column (4) of Table 3. Moving along the horizontal axis, the figure indicates increases in the probability of working a small number of days per week, larger than the increase in the probability of working 0 days. Thus, the employment guarantee shifts the distribution of days worked to the left. This shift results in some women exiting the labor force, while many more reduce their annual days of work, reinforcing the negative impact of the employment guarantee on the earnings women bring into their households.


Panels (b) and (d) of Figure 5 are analogous in format to Panels (a) and (c) for rural married men. Panel (b) indicates an average treatment-control difference in the annual take-up of employment-guarantee jobs of 4.2, which is also outweighed by a negative impact on the average annual days worked in other activities. Consistent with husbands being primary workers, the employment guarantee does not affect their extensive-margin participation. However, it reduces their average annual days worked. Reallocation across activities explains the simultaneous reduction in days worked and the massive aggregate provision of employment-guarantee jobs for both men and women. An economic explanation is next.

5.1 The Employment Guarantee as Insurance of Household Earnings

In brief, the explanation is the following: the employment guarantee insures household earnings against shocks that make their available working days irregular. If these shocks are positively correlated, the insurance shuts down a motive for precautionary savings, thus increasing household consumption of goods and time spent in non-market activities (i.e., it increases time spent outside of work for both women and their husbands). Cultural norms are such that the impact on women affects their labor force participation along the extensive

²⁶The baseline refers to the 2004-2005 period, prior to the employment guarantee, in treatment states. For rural married men, the baseline averages of annual days worked in self-employment, agricultural wage labor, and non-agricultural wage labor are 77.3, 57.3, and 83.9, respectively.

Figure 5. Annual Days Worked and the Employment Guarantee



Note: Panels (a) and (c) display estimates of the aggregate average treatment on the treated for rural married women based on Equation (5) for each of the dependent variables labeled in the horizontal axes. Days worked are measured annually. Individuals who do not work in a certain category are assigned 0 days (i.e., days worked are not conditional on participation). The specification of Equation (5) includes individual, age, and spouse-age fixed effects, as well as the district-level and state-level controls in Table 2 (entered into the equation linearly). The confidence intervals are based on the jackknifed wild-bootstrapped distribution clustered at the state level. Panels (b) and (d) are analogous in format to Panels (a) and (c) for rural married men. **Sample:** Rural married female (a and c) and male (b and d) subsamples of the IHDS labor-market sample.

margin, not only along the intensive margin as is the case for men. I elaborate on this mechanism and the empirical implications for consumption and savings next.

I assume a household composed of a woman and her husband derives utility from the consumption of goods, C , and the total days spent in non-market activities, H .²⁷ The couple has a total of 730 days per year (365 per person), allocated between H and their total annual days worked, D . Economic uncertainty (risk) may limit their work days. If risk averse, they accumulate a buffer stock (precautionary savings) in anticipation of such uncertainty. In the absence of constraints, their optimal choice is the bundle (C^*, H^*) . With constraints, they need to work more than $D^* := 730 - H^*$ to accumulate a buffer stock (whenever these additional days are available). If their utility function is concave in both C and H , adjustments to accumulate the buffer stock involve reducing both C and H to moderate deviations from (C^*, H^*) .

In the absence of the employment guarantee, households face irregular restrictions and thus increase their total annual days worked, D , in anticipation of future (stochastic) restrictions. This includes both the primary (men) and secondary (women) workers, the latter of whom might not work absent these restrictions. With the employment guarantee, the risk of failing to meet D^* is mitigated,²⁸ removing the need to accumulate a buffer stock. Figure 5 supports this implication, suggesting that shocked households take up employment-guarantee jobs to reach D^* , while non-shocked households reduce their overall workdays, moving from buffer stock accumulation levels toward D^* .

An implication of the employment guarantee as insurance of household earnings is a decrease in the variance of consumption across the life cycle of the household members. If shocks are positively correlated, insurance not only decreases the variance but also increases average consumption (Blundell et al., 2016; Meghir and Pistaferri, 2011): those not shocked can increase their consumption in consecutive periods without needing to put money aside as a buffer, and, thus, the average increases. This response should also decrease average household savings.²⁹

Tests on the implications for household consumption and savings are in Figure 6 and Appendix Table A.6. Figure 6 illustrates the raw event study based on Equation (1) and the

²⁷This assumption is justified if the utility from such activities outweighs their cost. For instance, a woman might enjoy raising her children because it aligns with social norms or provides personal fulfillment, despite its costs. If the net benefit is strictly concave, the assumption about non-market activities holds.

²⁸Should the days provided by the employment guarantee fall short, households would still reduce their workdays and increase consumption, moving closer to the optimal bundle (C^*, H^*) .

²⁹The implication on savings is salient because the average of total days worked across activities decreases for both women and their husbands while the average daily wage remains unaltered (see Section 7).

subsample of rural households from the HE-NSS consumption sample, indicating an average increase of 5.3% in average monthly household consumption per capita from a baseline of 244.7 (2018 USD, PPP). Appendix Figures A.14 to A.16 confirm very similar results across other specifications based on Equation (2), while Appendix Table A.5 further corroborates these findings with longitudinal data on log household consumption per capita from the IHDS consumption sample. If household consumption is considered a metric of well-being, then the employment guarantee contributes to a reduction in household-level absolute poverty.

Household savings are only observed in the IHDS consumption sample. Appendix Table A.6 indicates that the employment guarantee decreases monthly household savings by 99.6 from a baseline average of 237.5 (2018 USD, PPP). The impact is large but unsurprising, as the employment guarantee represents a permanent source of insurance. Appendix Table A.6 also indicates a small but precisely estimated decrease in the likelihood of livestock ownership. A larger decrease in the more liquid form of assets is expected. The evidence on consumption and savings supports the interpretation of the employment guarantee as insurance of household earnings; it rationalizes the findings in Section 4.


6. Intra-Household Resource Sharing and the Employment Guarantee

I use a collective model of household decisions (e.g., Chiappori, 1988, 1992) to quantify the within-household distributional consequences of the employment guarantee. Such quantification requires imposing additional structure on the household decision-making process. I assume that the framework discussed in Section 5.1 represents the first stage of this process. In this stage, the woman and her husband decide on their annual total household consumption (of market goods), C . In the second stage, they decide how to allocate this total into different goods. An assumption is required for the second stage to be informative regarding the overall distribution of resources within the household. Namely, the structural parameters dictating the within-household distribution of resources in the second stage summarize such distribution in the general household problem described by the two stages. I model the second stage as follows.

Allocation of Total Household Consumption. An individual can be one of two types: woman (w) or husband (h). As before, I index the model elements by time relative to the start of the employment guarantee: $g \in \{\text{before}, \text{after}\}$ and treatment status: $d \in \{\text{control}, \text{treatment}\}$. These two indices define four regimes. The household allocates total consumption of goods, C_g^d , by solving

$$\max_{\mathbf{z}_g^d} \tilde{U}_g^d [U_g^{d,w}(\mathbf{x}_g^{d,w}), U_g^{d,h}(\mathbf{x}_g^{d,h})] \quad (6)$$

Figure 6. (Raw) Log of Household Consumption per Capita of Rural Households and the Employment Guarantee, Main Event Study



Note: Panel (a) displays estimates of γ_g^p for each quarter g and phase p based on Equation (1) using log household consumption as the dependent variable and without using either fixed effects or controls (i.e., raw). The estimates displayed are the quarterly (conditional) average in Phase- p districts located in treatment states minus the analogous average in Phase- p districts located in control states. These treatment-control differences are relative to the treatment-control difference in the closest period to 0 before implementation (reference period). The treatment-control difference in the reference period is thus set to 0 and appears in the plot without a confidence interval. Panel (b) displays the population-weighted average of the γ_g^p estimates in Panel (a) based on Equation (3). Both panels display the 95% confidence interval based on the jackknifed wild-bootstrapped distribution clustered at the state level for each treatment-control difference. Both panels display the average weighted treatment-control difference across phases before and after implementation, relative to the reference period, based on Equation (4). The jackknifed wild bootstrapped p -value clustered at the state level associated with the null hypothesis of 0 accompanies each of these differences. **Sample:** Rural subsample of the HE-NSS consumption sample.

subject to

$$\begin{aligned} \text{total household consumption of goods}_g^d &=: C_g^d = \mathbf{p}_g^d \cdot \mathbf{z}_g^d \\ \mathbf{z}_g^d &= \mathbf{A}_g^d [\mathbf{x}_g^{d,w} + \mathbf{x}_g^{d,h}], \end{aligned}$$

where $\tilde{U}_g^d$ is the (strictly concave) household utility function over consumption goods and $U_g^{d,r}(\cdot)$ is the corresponding (strictly concave) individual utility function of type $r \in \{w, h\}$.

Given C_g^d , the woman and her husband maximize household utility by buying a bundle of goods $\mathbf{z}_g^d$ at price $\mathbf{p}_g^d$ in the market. The block-diagonal matrix $\mathbf{A}_g^d$ characterizes a Gorman (1976) linear technology describing how the consumption of each item in the vector $\mathbf{z}_g^d$ is shared between them. The vector $\mathbf{x}_g^{d,r}$ is what individual r actually consumes. If an element of the diagonal of $\mathbf{A}_g^d$ is greater than 1, the corresponding good in $\mathbf{z}_g^d$ is shared. In this case, the purchased good is less than the sum consumed by w and h . If an element of the diagonal is 1, there is no sharing of the corresponding good. Put differently, sharing results in consumption of greater value than the nominal value of what the household purchases in the market.³⁰ Let $\tilde{\mathbf{p}}_g^d$ denote the (shadow) price, which adjusts $\mathbf{p}_g^d$ for the gains of sharing. If at least one good is shared, $\tilde{\mathbf{p}}_g^d \leq \mathbf{p}_g^d$. Dunbar et al. (2013) show that $\tilde{\mathbf{p}}_g^d = \mathbf{A}_g^d \mathbf{p}_g^d$ in this allocation problem.

Intra-Household Share of Resources. The allocation problem is Pareto efficient, which does not mean that the resulting optimal allocation is balanced between the woman and her husband or that the woman has a high bargaining power. The contract curve may contain a point where most expenditure is allocated towards $\mathbf{x}_g^{d,h}$ and away from $\mathbf{x}_g^{d,w}$. The Pareto weight is the marginal change in $\tilde{U}_g^d$ due to a unit increase in $U_g^{d,w}(\cdot)$; it summarizes her bargaining power relative to that of her husband and has a one-to-one relationship with the female intra-household share of total resources. I denote this share by $\eta_g^{d,w}$ (the corresponding husband share is $\eta_g^{d,h} := 1 - \eta_g^{d,w}$). Identifying and estimating the impact of the employment guarantee on $\eta_g^{d,w}$ allows me to quantify the within-household distributional consequences of this policy. The identification challenge is that, usually, $\mathbf{z}_g^d$ is observed while $\mathbf{A}_g^d$, $\mathbf{x}_g^{d,w}$, and $\mathbf{x}_g^{d,h}$ are not, making direct computation of $\eta_g^{d,w}$ impossible.

An Engel-Curve System for Assignable Private Goods. Private assignable goods are goods for which (i) the relevant diagonal entry of $\mathbf{A}_g^d$ equals 1; and (ii) the analyst can assign them to either the woman or her husband. They are helpful for identification because,

³⁰Suppose the market price of sandwich units is 1. A woman and her husband want to consume 4 units of sandwich each. The relevant entry of $\mathbf{x}_g^{d,w} + \mathbf{x}_g^{d,h}$ is 8. If sharing inputs allows them to save 20% of the preparation cost, the relevant entry of $\mathbf{A}_g^d$ is 1.2. The market value of their sandwich units is $8 \cdot 1.2 = 10$, which is the relevant entry of $\mathbf{z}_g^d$.

for these goods, the market and shadow prices are the same, bypassing the fact that $\mathbf{A}_g^d$ is not observed. In practice, I observe the composites of private assignable goods for women and husbands described in Section 2.2. Dunbar et al. (2013) show that, without additional assumptions, the Engel curve for the composite of $r \in \{w, h\}$ is

$$\Omega_g^{d,r}(C_g^d) = \eta_g^{d,r}(C_g^d) \cdot \omega_g^{d,r}(\eta_g^{d,r}(C_g^d) \cdot C_g^d) \quad (7)$$

for a given price vector $\mathbf{p}_g^d$.³¹ $\Omega_g^{d,r}(C_g^d)$ is the share of total household consumption devoted to the private assignable composite of $r \in \{w, h\}$. $\omega_g^{d,r}(\eta_g^{d,r}(C_g^d) \cdot C_g^d)$ is the share of total consumption that an individual of type r would devote to their composite of private assignable goods if they were to allocate $\eta_g^{d,r} \cdot C_g^d$ to maximize $U_g^{d,w}$ by buying goods $\mathbf{x}_g^{d,w}$ at prices $\tilde{\mathbf{p}}_g^d$ (i.e., it is the Engel curve of the decentralized problem).³²

Identification. Dunbar et al. (2013) propose the following identification argument. Suppose that (i) the share $\eta_g^{d,r}$ is independent of the level of total household consumption; and (ii) the Engel curve is log-linear. Then, $\omega_g^{d,r}(C_g^d) = \alpha_g^{d,r} + \beta_g^{d,r} \cdot \log(C_g^d)$. Assumption (i) is an exclusion restriction. It states that $\eta_g^{d,r}$ is independent of C_g^d but not that $\omega_g^{d,r}$ is independent of C_g^d . It is a plausible assumption for describing a relatively homogenous population. Additionally, in my empirical strategy, I estimate the Engel curves for each regime, allowing some heterogeneity. Assumption (ii) is a shape restriction. If it did not hold, $\omega_g^{d,r}$ would still only be a function of C_g^d for a given price vector $\mathbf{p}_g^d$. However, the relationship would not be log-linear. Examples of demand systems where the relationship is log-linear include the “almost ideal” demand system (Deaton and Muellbauer, 1980). Assumption (iii), an additional shape restriction, is $\beta_g^{d,w} = \beta_g^{d,h} =: \beta_g^d$. Assumptions (i), (ii), and (iii) allow me to rewrite Equation (7) as:

$$\Omega_g^{d,r}(C_g^d) = \underbrace{a_g^{d,r}}_{\text{constant}} + \underbrace{b_g^{d,r}}_{\text{slope}} \cdot \log(C_g^d) + \xi_g^{d,r}, \quad (8)$$

where $a_g^{d,r} := \eta_g^{d,r} \cdot (\alpha_g^{d,r} + \beta_g^d)$, $b_g^{d,r} := \eta_g^{d,r} \cdot \beta_g^d$, and $\xi_g^{d,r}$ is an error term. The additional shape restriction indicates that differences between the woman and her husband in the share spent in private assignable consumption are summarized by $\alpha_g^{d,r}$ and not $b_g^{d,r}$. Lechene et al. (2022) show that recasting the Engel curves as a function of $a_g^{d,r}$ and $b_g^{d,r}$ allows identifying $\eta_g^{d,r}$ by noting that, under the three assumptions, $\eta_g^{d,w} = b_g^{d,w} / (b_g^{d,w} + b_g^{d,h})$, which provides a

³¹I suppress prices because I do not rely on them for identification. This is an advantage, as identification of resource shares in collective models sometimes requires price variation (Chiappori and Mazzocco, 2017).

³²For example, suppose that, for a given value of C_g^d , $\eta_g^{d,w} = 0.5$. If a woman spends 10% of total consumption in her private assignable composite in the decentralized problem, the share of total household consumption spent in this composite in the household problem is $0.5 \cdot 0.10 = 0.05$.

plug-in estimator of $\eta_g^{d,w}$ once estimates of $b_g^{d,w}$ and $b_g^{d,h}$ are available.

Though identification of $\eta_g^{d,w}$ relies on exclusion and shape restrictions, it has a transparent “reduced-form” interpretation. Namely, it is the response of the share of the private assignable composite good of the woman ($\Omega_g^{d,w}$) when total household consumption (C_g^d) increases. The response ($b_g^{d,w}$) is relative to the total of the responses of the woman and her husband ($b_g^{d,w} + b_g^{d,h}$). The larger the woman’s response relative to the overall household response, the larger her share $\eta_g^{d,w}$ and, thus, her bargaining power. An example relevant to this context is the following. Suppose that, as household consumption C_g^d increases, $\Omega_g^{d,m}$ and $\Omega_g^{d,w}$ increase. If $|b_g^{d,h}| > |b_g^{d,w}|$, the husband is able to increase consumption of his composite more than the increase of the woman. Identification is only achieved if $|\beta_g^d| = |(b_g^{d,w} + b_g^{d,h})| \neq 0$. Lechene et al. (2022) propose testing whether this condition holds by estimating the sum of Equation (8) across w and h :

$$\begin{aligned} \Omega_g^d(C_g^d) &=: \Omega_g^{d,w}(C_g^d) + \Omega_g^{d,h}(C_g^d) \\ &=: \underbrace{\alpha_g^d}_{\text{constant}} + \underbrace{\beta_g^d}_{\text{slope}} \cdot \log(C_g^d) + \underbrace{\xi_g^d}_{:=\xi_g^{d,w} + \xi_g^{d,h}}. \end{aligned} \quad (9)$$

Estimation. Ordinary least-square estimation of Equation (8) for $r \in \{w, h\}$ and Equation (9) yields unbiased estimates of $b_g^{d,w}$, $b_g^{d,h}$, and β_g^d under the standard mean-independence assumption on $\xi_g^{d,r}$ for $r \in \{w, h\}$. If C_g^d is measured with error, this assumption does not hold. I thus estimate the three equations instrumenting $\log(C_g^d)$. The estimation uses the rural subsample of the HE-NSS household consumption sample, where I observe the shares of the private assignable private composite goods for $r \in \{w, h\}$, $\log(C_g^d)$, and an alternative measure of log total household consumption, which I use as an instrument for $\log(C_g^d)$.³³ With the estimates of $b_g^{d,w}$, $b_g^{d,h}$, and β_g^d , I provide inference on the identification test $\beta_g^d = 0$ for each regime. I also estimate $\eta_g^{d,w}$ and the ATT on this parameter based on the expression $\left[\left(\eta_{\text{after}}^{\text{treatment},w} - \eta_{\text{after}}^{\text{control},w} \right) - \left(\eta_{\text{before}}^{\text{treatment},w} - \eta_{\text{before}}^{\text{control},w} \right) \right]$.³⁴

³³The consumption measures in the HE-NSS are reported monthly. In Section 5, I use monthly total household consumption resulting from adding itemized consumption of all observed goods. The alternative measure used as an instrument is based on a variable directly measuring overall annual total household consumption. Note that this strategy tackles measurement error but does not tackle more general concerns related to endogeneity.

³⁴The findings in Section 5.1 show an increase in household consumption, leading to a movement along the Engel curve. For the ATT on the share parameter to reflect this movement and be informative about female well-being, the Engel curve must remain structurally invariant: the utility function parameters that summarize preferences of women and their husbands should not change across different policy regimes. If this condition is met, then movements along the Engel curve generate the treatment effect on the share parameter. The methodology here does not permit a comprehensive test of invariance. However, note that Panel d. of

Structural Parameters and the Employment-Guarantee Impact On Them. Panels a. and b. of Table 4 summarize the assignable private good composites for women and husbands. They display the monthly averages in 2018 PPP dollars. For reference, the averages of monthly total household consumption for treatment and control states before the employment guarantee are 1,293 and 1,201 (2018 USD, PPP). The average expenditure in the husband composite increases more than the average expenditure in the woman’s in treatment states after netting out the control-group after-before difference. Preliminarily, these averages suggest a negative ATT on $\eta_g^{d,w}$. However, recall that $\eta_g^{d,w}$ is the slope of the Engel curve of the female private assignable good in relative terms. The impact of the employment guarantee in such a relative slope drives the ATT on the parameter.

Panel c. displays baseline estimates of $\eta_g^{d,w}$. These are estimates of the female share of intra-household consumption in control and treatment states before the employment guarantee. They range between 0.42 and 0.45. This range is consistent with recent estimates in the literature. For example, Calvi (2020) obtains an estimate of 0.44 when using a nationally representative cross-section of households in India. The panel also displays an estimate of the ATT on $\eta_g^{d,w}$, which indicates that the employment guarantee decreases the female intra-household share of resources by 0.04 (s.e. 0.01). That is, it decreases the female share by 9% from the treatment-state baseline. This impact more than doubles the pre-treatment gap between the intra-household share of resources of the woman and her husband. It implies that, within the household, female bargaining power decreases as a result of the employment guarantee.³⁵

The structural results suggest that while the employment guarantee benefits households as a whole, it hurts the women within them. Precisely, it crowds out the labor force participation of rural married women. This crowd-out reduces their command of household earnings and, thereby, their share of intra-household resources. The structural evidence is necessarily based on untestable assumptions regarding the household decision-making process. Aiming to consolidate this evidence, I quantify its implications next.

A Measure of Domestic Independence within the Household. The negative impact

Table 4 indicates invariance in $\beta_g^d = (b_g^{d,w} + b_g^{d,h})$ across treatment regimes before the implementation of the employment guarantee, making it plausible that preference parameters remain invariant upon implementation of the employment guarantee.

³⁵Panels d. and e. of Table 4 provide identification tests. That is, estimates of β_g^d and inference on them and the standard first-stage instrumental-variable rank test for the estimation of Equation (8) for $r \in \{w, h\}$ and Equation (9). These tests are satisfactory for the four regimes. Note that the F statistics are very large, which is expected given that the two measures of total household consumption closely track each other. Note also that, in this instrumental-variable strategy, estimation of Equation (8) for $r \in \{w, h\}$ and Equation (9) has the same first stage.

Table 4. Summary of Structural Estimation of Household Resource-Allocation Model

(1)	(2)	(3)	(4)
Panel a. Average Female Private Assignable Composite (2018 USD, PPP)			
Before, Control	Before, Treatment	After, Control	After, Treatment
24.085	22.438	26.771	27.250
Panel b. Average Husband Private Assignable Composite (2018 USD, PPP)			
Before, Control	Before, Treatment	After, Control	After, Treatment
38.759	43.578	41.378	51.885
Panel c. Structural Parameter: Female Share of Intra-Household Resources			
$\eta_{\text{before}}^{\text{control}}$	$\eta_{\text{before}}^{\text{treat}}$	ATT	N
0.419 (0.006)	0.450 (0.018)	-0.040 (0.011)	84,988
Panel d. Identification (Rank) Test: Engel-Curve Slope			
Before, Control	Before, Treatment	After, Control	After, Treatment
0.023 (0.007)	0.022 (0.008)	0.019 (0.006)	0.018 (0.008)
Panel e. Identification (Rank) Test: First Stage F-stat			
Before, Control	Before, Treatment	After, Control	After, Treatment
13,328.1	35,633.0	6,315.2	2,267.0

Note: Panel a. and b. summarize the average expenditure in the female and husband private assignable composite goods for treatment and control states before (2004-2005) and after (2011-2012) the implementation of the employment guarantee. The units are monthly expenditure in 2018 USD, PPP. Panel c. displays estimates of the female intra-household share of resources in the control ($\eta_{\text{before}}^{\text{control}}$) and treatment ($\eta_{\text{before}}^{\text{treat}}$) states before implementation, as well as an estimate of the aggregate average treatment on the treated on the share and the number of observations (**N**). Panels d. and e. provide the rank tests for the system of equations identifying the female intra-household share of resources. Panel d. displays estimates of the slope (β_g^d) in Equation (9) for treatment and control states before implementation. Panel e. displays the corresponding F statistics from the first stage in which I instrument log total household consumption with its alternative measure when estimating Equation (8) for $r \in \{w, h\}$ and Equation (9). In Panels c. and d., the standard errors (in parentheses) are based on the jackknifed wild-bootstrapped distribution clustered at the state level. The F statistics are asymptotic and clustered at the state level.

Sample: Rural subsample of the HE-NSS consumption sample, limited to the observations in rounds 61 (2004-2005) and 68 (2011-2012).

of the employment guarantee on the female intra-household share of resources likely decreases their economic independence. Such a decrease could deteriorate the relationship between women and their husbands (Anderson, 2021), increasing domestic abuse and intimate-partner violence.³⁶ I test whether the employment guarantee generates this deterioration using the

³⁶This argument directly links the decrease in female labor force participation, bargaining power, and intra-household share of resources to the decrease in domestic independence—the different pieces in this section point towards this link. However, the decrease in female labor force participation could also directly decrease domestic independence (even if bargaining power and intra-household share were fixed). For example, if women and husbands disagreed about time allocation after the decrease in female labor force

DiD estimator in Equation (5) and longitudinal data on “domestic independence” available in the IHDS female well-being sample. Panel a. of Table 4 displays the impact on this measure.³⁷ It indicates that the employment guarantee limits domestic independence, decreasing the index by 0.313—the variable is standardized to an in-sample mean of 0 and a standard deviation of 1. The basic falsification tests using the placebo samples of rural married non-disadvantaged and urban married women indicate a small impact that does not differ statistically from 0 at standard significance levels.³⁸

BMI: A Well-Being Measure. Longitudinal data on BMI, also available in the IHDS female well-being sample, allows me to corroborate the result based on domestic-independence measures. This corroboration is important because the domestic-independence measures are inherently subjective and sensitive, and, thus, prone to measurement error. These measures also have a high item non-response rate. I use BMI because it has been linked to the intra-household distribution of resources in India (Anderson and Genicot, 2015; Calvi, 2020), and, more generally, it has been used as a measure of overall mental and physical well-being.

Panel b. of Table 5 indicates that the employment guarantee reduces BMI by 0.35 points. The basic placebo tests in the falsification samples are satisfactory. Data on height allow me to provide further corroboration. There is no reason for the employment guarantee to have an effect on height because the youngest women in the sample are 24 years old, which is after the typical age at which Indian women stop growing (Khadilkar et al., 2009). Panel c. of Table 5 verifies that, while it differs from 0 statistically when using standard significance levels, the impact on height for rural married women is small. The negligible impact on height confirms that the impact on BMI for rural married women is driven by an effect on weight. For the adult women with low baseline weight and stable height that I analyze, a loss of BMI increases the risk of all-cause mortality (Thorogood et al., 2003).³⁹

The evidence presented in this section has limitations, relying on exclusion and shape restrictions for identifying and estimating structural parameters, or on the IHDS female well-

participation and such disagreement generated violence.

³⁷For the results in Table 5, I include individual, age, and spouse-age fixed effects as well as the district-level and state-level controls in Table 2 (entered into the equation linearly) in the estimation of Equation (5). Appendix Tables A.7 to A.9, based on alternative specifications, show very similar results.

³⁸In Table 5, I use all non-disadvantaged married women (rural and urban) as opposed to only those who are rural as in the rest of the paper. This allows me to increase the size of the subsample, which is necessary due to the more limited number of observations in the female well-being sample of the IHDS.

³⁹There is a U-shaped relationship between physical and mental health and BMI (both low and high levels of BMI are detrimental individual health; Allison et al., 1997; de Wit et al., 2009). A decrease in BMI from the low baseline average of 20.9 in treatment states makes women vulnerable physically and mentally (Selvamani and Singh, 2018). A “healthy” decrease in BMI due to a reduction from a high baseline value is not salient in the sample I analyze (only about eleven percent of women are overweight or obese).

Table 5. Female Well-Being and the Employment Guarantee

	(1)	(2)	(3)
	<u><i>Independence Index</i></u>	<u><i>Body-Mass Index</i></u>	<u><i>Height (Meters)</i></u>
Panel a. Rural Married Women			
Estimate	-0.313	-0.350	0.012
(<i>p</i> -value)	(0.025)	(0.005)	(0.007)
Baseline Treatment Mean	0.114	20.853	1.515
Observations	5,962	8,722	8,722
Panel b. Non-Disadvantaged Married Women			
Estimate	-0.184	-0.043	-0.005
(<i>p</i> -value)	(0.491)	(0.894)	(0.442)
Baseline Treatment Mean	0.289	22.332	1.531
Observations	2,626	4,021	4,021
Panel c. Urban Married Women			
Estimate	-0.111	-0.013	0.003
(<i>p</i> -value)	(0.725)	(0.951)	(0.793)
Baseline Treatment Mean	0.336	22.812	1.523
Observations	3,114	4,188	4,188

Note: Panel a. displays details from the estimation of the aggregate average treatment on the treated for rural married women based on Equation (5) using the dependent variable indicated in the column, including individual, age, and spouse-age fixed effects and the district-level and state-level controls in Table 2 (entered into the equation linearly). Panels b. and c. are analogous in format to Panel a. for non-disadvantaged and urban married women. For each estimate, the state-clustered jackknifed wild-bootstrapped *p*-value associated with the null hypothesis of 0 is displayed in parentheses. **Sample:** Subsamples of the IHDS female well-being sample indicated in the label.

being sample, which suffers from a large item non-response rate. Despite these imperfections, the findings are cohesive across diverse sources. While the employment guarantee generally benefits households, it adversely affects the women within them. A joint interpretation of the results and those from Sections 4 and 5 indicates that the employment guarantee crowds out the labor force participation of rural married women, reducing their command of household earnings and, thus, their intra-household share of consumption and overall well-being. A final qualitative exploration of the impact on labor force participation solidifies this narrative.

Further Linking the Impact on Women to Cultural Norms. Eswaran et al. (2013) argue that “the time allocation of married women to market work, especially in rural areas, is mediated by their family’s desire to maintain «status»” and that “working outside the home is deemed to be a low-status activity for married, rural women.” They argue that, as economic conditions improve, “status concerns become more salient and married women may gradually begin to withdraw from market work.” In this setting, socioeconomic status also determines the targeting of the employment guarantee. Panel (a) of Figure 7 further illustrates that the impact is driven by those who are disadvantaged using the ATT estimator in Equation (5) and the IHDS labor-market sample. The employment guarantee improves the economic conditions of those in disadvantaged households, and, as a response, they become more likely to quit the labor force.

Eswaran et al. (2013) take the argument further. They propose that social norms dictating family preference for women not to work originate from patriarchal regimes, “where contact with males outside the household was deemed a «polluting» influence that was to be avoided where possible.” If their argument held, the more patriarchal the state, the larger the observed negative impact on female labor force participation should be. I classify treatment states by a patriarchy index provided in Singh et al. (2022). I group them into three groups of two states, from least patriarchal to most patriarchal. I then estimate the ATT on the rural married disadvantaged for each of these three groups as the treatment group and the control group of states used throughout the paper.⁴⁰ Panel (b) of Figure 7 shows that the more patriarchal the state group, the likelier the disadvantaged women in Panel (a) are to drop from the labor force as a result of the employment guarantee. Figure 7 further suggests that cultural norms are essential in determining the labor-market crowd-

⁴⁰Rajasthan is excluded from this state-wise comparison as it is the only state with a positive treatment effect on rural female labor force participation, attributable to a history of effective casual-job programs (see Section 7). Including Rajasthan complicates the analysis due to its unique context compared to other states lacking similar programs. This approach aligns with other studies (e.g., Azam, 2011) and is supported by Appendix Table A.10, which confirms that excluding Rajasthan does not undermine the robustness of the results. Rajasthan is included in all other analyses in the paper, where I aim to understand the national implications of the employment guarantee, as opposed to state-level specificities.

out of rural married women and, thus, its within-household distributional consequences. It indicates that, in states with a stronger patriarchal history, the employment guarantee further damages female well-being.

7. Comparison to Other Studies


While some studies of the employment guarantee use the EU-NSS, they do not use all of its available rounds in combination. I use seven rounds of the EU-NSS. Other studies use two rounds (e.g., Azam, 2011; Imbert and Papp, 2015; Misra, 2019), one round before the start of the employment guarantee (2004-2005) and one round after (2007-2008). Using seven rounds allows me to study longer-term impacts and verify the absence of trends in several periods before implementation. It also allows me to use event-study methods developed for evaluating programs with staggered roll-out (see de Chaisemartin and d’Haultfoeuille, 2020). Previous studies focus on one data source. I combine the EU-NSS, the HE-NSS, and the IHDS, corroborating findings based on different datasets and empirical strategies and testing implications of my theoretical framework on a variety of outcomes.

Appendix 7 provides an empirical comparison to Azam (2011). This comparison is relevant in itself because his findings appear to contradict mine. It is also relevant because succeeding studies use the same or very similar strategies (e.g., Imbert and Papp, 2015; Misra, 2019). Azam (2011) finds that the employment guarantee increases female labor force participation by 2.4 percentage points. He relies on the district-level variation in treatment timing and does not use the state-level variation in treatment intensity.

Appendix 7 indicates that the difference between the results in Azam (2011) and this paper is unlikely to be driven by sample composition (e.g., a specific age profile, marital status, or state of residence) or specification of controls. The difference is more likely to be driven by the focus on different parameters of interest. I focus on a longer-term impact. I observe individuals up to five years after the employment guarantee. Azam (2011) focuses on a short-term impact. He observes individuals at most two years after the employment guarantee. That is also the case in other studies (e.g., Imbert and Papp, 2015; Misra, 2019). His strategy uses the majority of rural Indian districts, those in Phase 3, as the control group. By construction, his parameter estimates do not contain the impact on Phase-3 districts or the longer-term dynamics driving treatment effects.

Appendix 7 also analyzes the impact on rural wages. This analysis is relevant for three reasons. First, Sukhtankar (2016) indicates a positive impact on rural wages as a common finding in the literature. Second, Imbert and Papp (2015) find a positive impact on casual-work wages, which is part of the consensus documented in Sukhtankar (2016). Third,

Figure 7. Labor Force Participation of Rural Married Women and the Employment Guarantee, Disadvantage and Patriarchy



Note: Panel (a) displays estimates of the aggregate average treatment on the treated for rural married women by socioeconomic disadvantage, as defined in Table 2, using labor force participation as the dependent variable. The estimation is based on Equation (5), including individual, age, and spouse-age fixed effects, as well as the district-level and state-level controls in Table 2 (entered into the equation linearly). The confidence intervals are based on the jackknifed wild-bootstrapped distribution clustered at the state level. Panel (b) is analogous in format to Panel (a), displaying the estimate of the ATT using the states in the horizontal axis as the treatment group and the control states used throughout the paper as the control group. Panel (a) includes the state of Rajasthan in the treatment states, along with the treatment states labeled in Panel (b). **Sample:** Rural married female subsample of the IHDS labor-market sample.

determining the size and magnitude of such an impact is relevant for my economic interpretation of the employment guarantee in Section 5. Application of my empirical strategy yields a small negative average impact on rural daily wages across activities of -1.4% (s.e. 13.4%) from a baseline of 5.8 (2018 USD, PPP). This finding implies that the average positive impact of 4.7% (s.e. 2.3%) on casual-work wages reported by Imbert and Papp (2015) does not translate into a sizable impact on longer-term rural wages.⁴¹ My estimate is similar to the estimate of Zimmermann (2024), who, focusing on men between 18 and 60 years old and using a different identification strategy than mine, finds that the employment guarantee decreases rural wages across working activities by an average of -1.8% (s.e. 3.9%).

Imbert and Papp (2015) argue that, due to general-equilibrium effects, casual works competing with the employment guarantee increase their wages. I pursue an interpretation of the employment guarantee based on its direct impact as insurance of household earnings on household-level and individual-level decisions. I do so because, theoretically, it is difficult to sustain long-term wage changes as a main mechanism for the employment-guarantee impacts, especially without documented impacts on human capital. Additionally, the long-term impact on rural wages does not differ statistically from 0 when applying my empirical design. In contrast, I find that the direct impact is salient in magnitude and statistical significance for several implications on household and individual behavior.

Despite the differences with Azam (2011) and Imbert and Papp (2015), my findings broadly agree with other studies in terms of impacts on time allocation across working activities and household consumption. Table 6 summarizes a set of recent studies that generally coincide with the rest of the literature. Misra (2019) finds a reallocation of working activities towards public works (which include employment-guarantee jobs). Indeed, this reallocation is also documented by Imbert and Papp (2015) and Klonner and Oldiges (2022), who also find a decrease in agriculture as the main household occupation. Zimmermann (2024) finds a reallocation towards self-employment for men. All these reallocation results are consistent with Figure 5. The positive impacts on household consumption per capita documented by Bose (2017) and Klonner and Oldiges (2022) are also consistent with Section 5.

Sukhtankar (2016) states that research on the employment guarantee is still “badly needed,” but that “current standards for causal inference and the availability of data will remain high hurdles for those who wish to take on this challenge.” He states that the identification of mechanisms “demands even more from data and empirical methods.” This paper aims to fill some of the referenced gaps, relying on a strategy that examines longer-term

⁴¹Berg et al. (2018) and Klonner and Oldiges (2022) report findings on wages similar to Imbert and Papp (2015). Their strategies identify a short-term impact and are subject to the discussed selection caveats.

Table 6. Summary of Studies of the Employment Guarantee and Labor-Market and Consumption Outcomes

Source	Data Set	Years	Observation Units	Outcome	Policy Measure	Variation	Main Result
Azam (2011)	EU-NSS	2004-05, 2007-08	women ages 18-60	labor force participation	district-level guarantee presence	district-level rollout	labor force participation ↑
Bose (2017)	HE-NSS	2003, 2007-08	households in Phase-1 or Phase-3 districts of 19 major states	household consumption per capita	district-level guarantee presence	district-level rollout	consumption per capita ↑
Imbert and Papp (2015)	EU-NSS	2004-05, 2007-08	adults ages 18-60	casual-work wage/daily earnings	district-level guarantee presence	district-level rollout	casual-work wage ↓
Klonner and Oldiges (2022)	EU-NSS and HE-NSS	2007-08	households	household occupation and consumption	district-level guarantee presence	district-level rollout	agriculture main occupation ↓, consumption per capita ↑
Misra (2019)	EU-NSS	2004-05, 2007-08	adults ages 15-60	time in public/private works	district-level guarantee presence	district-level rollout	public works ↑, private works ↓
Zimmermann (2024)	EU-NSS	2007-08	men ages 18-60 in Phase-2 and Phase-3 districts	participation in work categories	district-level guarantee presence	district-level index cutoff for roll-out phase definition	private works ↓, self-employment ↑

Abbreviations: EU-NSS: Employment and Unemployment National Sample Survey. HE-NSS: Household Expenditure National Sample Survey. **Study Details:** Azam (2011): see in-text discussion for details. Standard errors are clustered at the district level. Bose (2017): Strategy is the same as in Azam (2011) but does not consider Phase-2 districts. Annual household consumption per capita increases 10.6% (s.e. 2.7%). Standard errors are clustered at the district level. Other outcomes analyzed: consumption categories (food and durable goods), education, and health. Imbert and Papp (2015): see in-text discussion for details. Standard errors are clustered at the district level. Klonner and Oldiges (2022): Strategy is regression discontinuity design based on index classifying districts into their implementation phases, thus making it possible to compare Phase-1 (early implementers) and Phase-2 (late implementers) districts at the eligibility threshold. Similarly with Phase-2 and Phase-3 district comparisons. Agriculture as main household occupation (reported in the Spring 2008 for treatment states, as classified in Section 3) decreases 13% (s.e. 4.2%). Household consumption per capita (reported in the Spring 2008 for treatment states, as classified in Section 3) increases 16% (s.e. 5.4%). Standard errors are clustered at the district level. Other details: authors supplement EU-NSS and HE-NSS with a survey of the NSS inquiring on education expenditure. Other outcomes analyzed: several. Misra (2019): Strategy is the same as in Azam (2011) but focuses on dry season. Further divides estimation by districts dominated and not dominated by landlord class. Main results are for districts not dominated by landlord class. Public works increase 0.936 pp. (s.e. 0.396) and private works decrease 2.927 pp. (s.e. 1.146). Standard errors are clustered at the district level. Other outcomes analyzed: wages. Zimmermann (2024): Strategy is the same as Klonner and Oldiges (2022) but focuses on Phase-2 and Phase-3 districts. Private employment decreases by 4.4 pp. (s.e. 2.6). Self-employment increases 4.9 pp. (s.e. 2.8). Standard errors are clustered at the district level. Other outcomes analyzed: wages.

impacts, documenting robustness across data sources, and examining economic mechanisms.

8. Summary and Final Comments

The labor force participation of women in India is salient for its low level and recent decrease, which contrasts with increasing trends around the world. Rural married women drove a recent countrywide decrease, lowering their participation by 25% from a baseline of 40%. This decrease mostly occurred between 2005 and 2012, and, thereafter, their participation remained around 30% through 2020. This paper aims to provide a reason for the low level and decrease. I argue that rural married women supply labor as added workers, i.e., only as a source of household insurance against economic uncertainties (risks) inherent to their households. An improvement in the economic conditions of their households thus increases their time spent in non-market activities. Social norms that establish family preferences for women not to work at all reinforce this potential increase. I find that this mechanism prevails upon the implementation of the Mahatma Gandhi National Rural Employment Guarantee Act. The employment guarantee effectively insures household earnings, replacing the role of women as added workers. While a fraction of married women take up employment-guarantee jobs, an even larger fraction reduce their participation in paid agricultural and non-agricultural activities as a consequence of the employment guarantee. The net impact on the labor force participation of rural married women is negative and explains up to 30% of the countrywide rural decrease. The insurance provided by the employment guarantee shuts down a motive for accumulating precautionary savings and increases average household consumption among rural households. Therefore, the employment guarantee reduces average poverty at the household level. However, by crowding out female labor force participation, the employment guarantee reduces the command of household earnings women have and, therefore, decreases their intra-household share of resources and overall well-being.

References

- Abadie, A., S. Athey, G. W. Imbens, and J. M. Wooldridge (2022). When should you adjust standard errors for clustering? *Quarterly Journal of Economics* 138(1), 1–35.
- Ackerson, L. K. and S. V. Subramanian (2008). Domestic violence and chronic malnutrition among women and children in India. *American Journal of Epidemiology* 167(10), 1188–1196.
- 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(3), 783–818.
- Afridi, F., A. Mukhopadhyay, and S. Sahoo (2016). Female labor force participation and

- child education in India: Evidence from the national rural employment guarantee scheme. *IZA Journal of Labor & Development* 5(1), 1–27.
- Alik-Lagrange, A. and M. Ravallion (2018). Workfare versus transfers in rural India. *World Development* 112, 244–258.
- Allison, D. B., M. S. Faith, M. Heo, and D. P. Kotler (1997). Hypothesis concerning the U-shaped relation between body mass index and mortality. *American Journal of Epidemiology* 146(4), 339–349.
- Anderson, S. (2021). Intimate partner violence and female property rights. *Nature Human Behaviour* 5(8), 1021–1026.
- Anderson, S. and M. Eswaran (2009). What determines female autonomy? Evidence from Bangladesh. *Journal of Development Economics* 90(2), 179–191.
- Anderson, S. and G. Genicot (2015). Suicide and property rights in India. *Journal of Development Economics* 114, 64–78.
- Attanasio, O. P. and V. Lechene (2014). Efficient responses to targeted cash transfers. *Journal of Political Economy* 122(1), 178–222.
- Azam, M. (2011). The impact of Indian job guarantee scheme on labor market outcomes: Evidence from a natural experiment. *SSRN Electronic Journal*.
- Bahal, G. (2019). Estimating the impact of welfare programs on agricultural output: Evidence from India. *American Journal of Agricultural Economics*.
- Banerjee, A., E. Duflo, C. Imbert, S. Mathew, and R. Pande (2020). E-governance, accountability, and leakage in public programs: Experimental evidence from a financial management reform in India. *American Economic Journal: Applied Economics* 12(4), 39–72.
- Berg, E., S. Bhattacharyya, D. Rajasekhar, and R. Manjula (2018). Can public works increase equilibrium wages? Evidence from India’s national rural employment guarantee. *World Development* 103, 239–254.
- Bhargava, S. (2018). Why did Indian female labor force participation decline?: Evidence from a model of household labor supply. *SSRN Electronic Journal*.
- Blumberg, R. L. and M. T. Coleman (1989). A theoretical look at the gender balance of power in the American couple. *Journal of Family Issues* 10(2), 225–250.
- Blundell, R., L. Pistaferri, and I. Saporta-Eksten (2016). Consumption inequality and family labor supply. *American Economic Review* 106(2), 387–435.
- Bose, N. (2017). Raising consumption through India’s national rural employment guarantee scheme. *World Development* 96, 245–263.

- Calvi, R. (2020). Why are older women missing in India? The age profile of bargaining power and poverty. *Journal of Political Economy* 128(7), 2453–2501.
- Cameron, A. C., J. B. Gelbach, and D. L. Miller (2008). Bootstrap-based improvements for inference with clustered errors. *Review of Economics and Statistics* 90(3), 414–427.
- Canay, I. A., A. Santos, and A. M. Shaikh (2021). The wild bootstrap with a “small” number of “large” clusters. *Review of Economics and Statistics* 103(2), 346–363.
- Chiappori, P.-A. (1988). Rational household labor supply. *Econometrica* 56(1), 63–90.
- Chiappori, P.-A. (1992). Collective labor supply and welfare. *Journal of Political Economy* 100(3), 437–467.
- Chiappori, P.-A. and M. Mazzocco (2017). Static and intertemporal household decisions. *Journal of Economic Literature* 55(3), 985–1045.
- Chief Labour Commissioner (2022). Minimum wages. Website, <https://clc.gov.in/clc/min-wages>.
- Chikermane, G. (2018). *70 Policies That Shaped India: 1947 to 2017, Independence to 2.5 Trillion*. Observer Research Foundation.
- de Chaisemartin, C. and X. d’Haultfoeulle (2020). Two-way fixed effects estimators with heterogeneous treatment effects. *American Economic Review* 110(9), 2964–96.
- de Wit, L. M., A. van Straten, M. van Herten, B. W. Penninx, and P. Cuijpers (2009). Depression and body mass index, a U-shaped association. *BMC Public Health* 9(1), 1–6.
- Dean, J. T. and S. Jayachandran (2019). Changing family attitudes to promote female employment. *AEA Papers and Proceedings* 109, 138–142.
- Deaton, A. and J. Muellbauer (1980). An almost ideal demand system. *American Economic Review* 70(3), 312–326.
- Desai, S. and O. Joshi (2019). The paradox of declining female work participation in an era of economic growth. *Indian Journal of Labour Economics* 62(1), 55–71.
- Dunbar, G. R., A. Lewbel, and K. Pendakur (2013). Children’s resources in collective households: Identification, estimation, and an application to child poverty in Malawi. *American Economic Review* 103(1), 438–471.
- Dutta, P., R. Murgai, M. Ravallion., and D. van de Walle (2012). Does India’s employment guarantee scheme guarantee employment? Technical Report WPS6003, World Bank.
- Dutta, P. V., R. Murgai, M. Ravallion, and D. Van De Walle (2014). Right to work?: Assessing India’s employment guarantee scheme in Bihar. Technical Report 85414, World Bank.

- Eswaran, M., B. Ramaswami, and W. Wadhwa (2013, January). Status, caste, and the time allocation of women in rural India. *Economic Development and Cultural Change* 61(2), 311–333.
- Fetzer, T. (2020). Can workfare programs moderate conflict? Evidence from India. *Journal of the European Economic Association* 18(6), 3337–3375.
- Field, E., R. Pande, N. Rigol, S. Schaner, and C. Troyer Moore (2021). On her own account: How strengthening women’s financial control impacts labor supply and gender norms. *American Economic Review* 111(7), 2342–75.
- Fletcher, E., R. Pande, and C. M. T. Moore (2017). Women and work in India: Descriptive evidence and a review of potential policies. *SSRN Electronic Journal*.
- Gorman, W. M. (1976). Tricks with utility functions. In *Essays in Economic Analysis*, pp. 211–243. Cambridge University Press.
- Government of India (2022). Open government data platform India. Website, <https://data.gov.in/>.
- Heath, R. and X. Tan (2020). Intrahousehold bargaining, female autonomy, and labor supply: Theory and evidence from India. *Journal of the European Economic Association* 18(4), 1928–1968.
- Heckman, J., N. Hohmann, J. Smith, and M. Khoo (2000). Substitution and dropout bias in social experiments: A study of an influential social experiment. *Quarterly Journal of Economics* 115(2), 651–694.
- Imbert, C. and J. Papp (2014). Estimating leakages in India’s employment guarantee: An update. Technical report, World Bank.
- Imbert, C. and J. Papp (2015). Labor market effects of social programs: Evidence from India’s employment guarantee. *American Economic Journal: Applied Economics* 7(2), 233–263.
- Imbert, C. and J. Papp (2019). Short-term migration, rural public works, and urban labor markets: Evidence from India. *Journal of the European Economic Association* 18(2), 927–963.
- Inter-University Consortium for Political and Social Research (2011). *India Human Development Survey Panel*. Ministry of Statistics and Programme Implementation: Government of India.
- Jayachandran, S. (2021). Social norms as a barrier to women’s employment in developing countries. *IMF Economic Review* 69(3), 576–595.
- Johnson, D. (2009). Can workfare serve as a substitute for weather insurance? The case of NREGA in Andhra Pradesh. *SSRN Electronic Journal*.

- Kessler-Harris, A. (2003). *In Pursuit of Equity: Women, Men, and the Quest for Economic Citizenship in 20th Century America*. Oxford University Press, USA.
- Khadilkar, V., A. Khadilkar, T. Cole, and M. Sayyad (2009). Cross-sectional growth curves for height, weight and body mass index for affluent Indian children, 2007. *Indian Pediatrics* 46, 477–489.
- Klasen, Janneke, S. P. (2015). What explains the stagnation of female labor force participation in urban India? *World Bank Economic Review* 29(3), 449–478.
- Klonner, S. and C. Oldiges (2022). The welfare effects of India’s rural employment guarantee. *Journal of Development Economics* 157, 102848.
- Lechene, V., K. Pendakur, and A. Wolf (2022). Ordinary least squares estimation of the intrahousehold distribution of expenditure. *Journal of Political Economy* 130(3), 681–731.
- Lundberg, S. (1985). The added worker effect. *Journal of Labor Economics* 3(1, Part 1), 11–37.
- Meghir, C. and L. Pistaferri (2011). Earnings, consumption and life-cycle choices. In *Handbook of Labor Economics*, Volume 4, pp. 773–854.
- Ministry of Rural Development (2005a). The National Rural Employment Guarantee Act.
- Ministry of Rural Development (2005b). The National Rural Employment Guarantee Act: Operational guidelines.
- Ministry of Rural Development (2007). The Maharashtra State Employment Guarantee Scheme, 2006. Website, <https://pib.gov.in/newsite/archiveReleases.aspx>.
- Ministry of Rural Development (2010). Districts under the Mahatma Gandhi National Employment Guarantee Act. Website, https://nrega.nic.in/MNREGA_Dist.pdf.
- Ministry of Statistics and Programme Implementation (2020a). National data archive: Employment and unemployment. Website, <http://microdata.gov.in/nada43/index.php/catalog/EUE/about>.
- Ministry of Statistics and Programme Implementation (2020b). National data archive: Household consumer expenditure. Website, <http://microdata.gov.in/nada43/index.php/catalog/CEXP/about>.
- Misra, K. (2019). Does historical land inequality attenuate the positive impact of India’s employment guarantee program? *World Development* 124, 1–22.
- Munshi, K. and M. Rosenzweig (2009). Why is mobility in India so low? Social insurance, inequality, and growth. Technical Report w14850, National Bureau of Economic Research.
- Office of the Registrar General and Census Commissioner (2020). Area and population. Website, https://censusindia.gov.in/Census_And_You/area_and_population.aspx.

- Planning Commission (2003). Report of the task force: Identification of districts for wage and self-employment programmes. Website, https://planningcommission.gov.in/aboutus/taskforce/inter/inter_reg.pdf.
- Qian, N. (2008). Missing women and the price of tea in China: The effect of sex-specific earnings on sex imbalance. *Quarterly Journal of Economics* 123(3), 1251–1285.
- Rahman, L. and V. Rao (2004). The determinants of gender equity in India: Examining dyson and moore’s thesis with new data. *Population and Development Review* 30(2), 239–268.
- Rangel, M. A. (2006). Alimony rights and intrahousehold allocation of resources: Evidence from Brazil. *Economic Journal* 116(513), 627–658.
- Selvamani, Y. and P. Singh (2018). Socioeconomic patterns of underweight and its association with self-rated health, cognition and quality of life among older adults in India. *PloS one* 13(3).
- Sen, A. (1990). Gender and cooperative conflicts. In A. Sen and I. Tinker (Eds.), *Persistent Inequalities*, pp. 235–297. Oxford University Press.
- Singh, A., P. Chokhandre, A. K. Singh, K. M. Barker, K. Kumar, L. McDougal, K. S. James, and A. Raj (2022, January). Development of the India patriarchy index: Validation and testing of temporal and spatial patterning. *Social Indicators Research* 159(1), 351–377.
- Sukhtankar, S. (2016). India’s national rural employment guarantee scheme: What do we really know about the world’s largest workfare program? In *India Policy Forum*, Volume 13, pp. 2009–10.
- Sun, L. and S. Abraham (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics* 225(2), 175–199.
- Thorogood, M., P. N. Appleby, T. J. Key, and J. Mann (2003). Relation between body-mass index and mortality in an unusually slim cohort. *Journal of Epidemiology and Community Health* 57(2), 130–133.
- Topalova, P. (2010). Factor immobility and regional impacts of trade liberalization: Evidence on poverty from India. *American Economic Journal: Applied Economics* 2(4), 1–41.
- World Bank (2022a). Female labor force participation rate. Website, <http://data.worldbank.org/indicator/SL.TLF.CACT.FE.ZS>.
- World Bank (2022b). Rural population. Website, <https://data.worldbank.org/indicator/SP.RUR.TOTL.ZS?locations=IN>.
- Zimmermann, L. (2024). Why guarantee employment? Evidence from a large Indian public-works program. *Economic Development and Cultural Change* 72(4), 2031–2067.