

Guaranteed Employment in Rural India: Richer Households, Poorer Women?

Jorge Luis García*

April 12, 2023

Abstract

I investigate the intra-household resource allocation consequences of a policy guaranteeing 100 annual days of minimum-wage employment to rural Indian households. The employment guarantee replaces married women as added workers, decreasing their labor force participation; it accounts for up to 30% of a 25-percent countrywide decrease in rural female labor force participation observed between 2005 and 2012 from a baseline of 40%. The employment guarantee insures household earnings and, thereby, increases household consumption. However, by crowding married women out of the labor force, it reduces their command of household earnings, intra-household share of consumption, and well-being.

JEL Codes: I31, I32, J12, J13, O12, O15.

Keywords: added-worker effect, family insurance, female labor force participation, guaranteed employment, intra-household bargaining power, poverty.

*John E. Walker Department of Economics, Clemson University. Address: 309-C Wilbur O. and Ann Powers Hall, Clemson, S.C. 29634. Telephone: (864)-656-3481. E-mail: jlgarci@clemson.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 Kabeer, 2008; 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 women decreased substantially during the last thirty years,¹ suggesting a worsening of their economic conditions. I investigate the intra-household resource allocation consequences of the Mahatma Gandhi Rural Employment Act and, thereby, document its potential contribution to the decreasing trend in the labor force participation of Indian women.

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. Precisely, its provision of jobs aims to insure households against the economic uncertainty generated by involuntary unemployment spells that are typical in 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 appears to be 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. During this period, the female-to-male labor force participation ratio decreased by 24% according to the same source. In contrast, this ratio increased by an average of 18% in the rest of South Asia. It also increased in the Middle East and North Africa (25%), Europe and Central Asia (11%), Latin America and the Caribbean (30%), and the countries belonging to the Organisation for Economic Co-operation and Development (17%).

household earnings, an employment-guarantee policy may compete with the role of women as “added” or “insurance” workers, which is common in rural India. 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, geo-identified 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 any positive impact on participation due to the provision of employment-guarantee jobs. I also find that the employment guarantee does not affect the participation rates of rural unmarried (never married, separated, divorced, or widowed) women or rural men of any marital status.

Most of the decrease in female labor force participation observed during the last thirty years occurred during the period observed in the EU-NSS sample. Precisely, 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 decrease for married women and similar percentages for all rural women.

I use the Indian Human Development Survey (IHDS) to construct a longitudinal, nationally representative, and geo-identified sample, including days worked by activity. In this sample, I estimate impacts relying on within-individual variation. I obtain an estimate of the impact on the labor force participation of rural married women identical to that obtained in the EU-NSS. In addition to this extensive-margin impact, I document an intensive-margin impact. The likelihood of working between zero and 90 days per year decreases by nine percentage points for these women. The employment guarantee thus shifts the distribution

of their days worked per year to the left and reinforces their role as added workers; it reduces their number of days worked per year by an average of 15 from a baseline average of 102. I also find that, while it does not impact rural married men on the extensive margin, the employment guarantee decreases their days worked by an average of 9.5 days from a baseline average of 217.

An economic framework explains the negative impact on days worked by rural married women and their husbands. Suppose that, together as a household, a woman and her husband derive strictly concave utility from bundles of consumption goods (household consumption) and days spent in non-market activities. Further, suppose that they are risk averse towards the days of work available to them, which are random due to negative shocks leading to involuntary unemployment. In this framework, a household accumulates a buffer stock in anticipation of negative shocks, financed 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 work days. This guarantee insures household earnings; it reduces the household's risk and thus its need to accumulate the buffer. While those households shocked in any given year take up employment-guarantee jobs, the average overall work days 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 days worked within my economic interpretation of the employment guarantee. 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; 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. They are likely to increase their non-market activities by reducing their participation in labor-market activities much more than their husbands. For some of them, such reduction includes completely quitting the labor force.

I also test the implication that the employment guarantee increases household consumption. Using cross-sectional and longitudinal data sources, I apply the same empirical design as when assessing labor-market outcomes. I document that the employment guaran-

tee *increases* monthly household consumption per-capita by an average of 6% to 8% from a baseline average of 256 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. This reduction is inconclusive regarding the within-household distributional consequences of the policy.

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 use longitudinal data from the IHDS to measure the impact of the employment guarantee on domestic independence, which I measure using a questionnaire asking women about views on domestic violence in their communities and whether they themselves need to ask permission from their husbands to perform activities such as visiting a friend or going to the store. I find that the employment guarantee decreases the domestic-independence 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) of women are also available in the IHDS. These data allow me to further corroborate the structural estimates. BMI is used to measure the consequences of changes in within-household resource allocation (e.g., Calvi, 2020); it strongly correlates with domestic violence and captures mental and physical health (Ackerson and Subramanian, 2008; Selvamani and Singh, 2018). I find that the employment guarantee has a substantial negative impact on BMI. The structural and reduced-form evidence indicate that, despite decreasing absolute household-level poverty, the employment guarantee makes women poorer within the household and hurts 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 a plausibly causal reason that is new to the literature. My interpretation of and evidence on why the employment guarantee crowds out female labor force participation and the intra-household consequences of this crowd-out are also new.

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) discusses social norms as a barrier to female employment in the developing world. 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. That is, 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, female health improves. Heath and Tan (2020) find that, as a consequence, women decide to participate more in the labor force.

The study by Field et al. (2021) helps to interpret my results and recommend policy 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 it reduces it 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 discussing 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.

Paper Plan. Section 2 describes the data analyzed. Section 3 describes the employment guarantee. Section 4 analyzes its impact on labor force participation. Section 5 explains the economic reasons for this impact and provides evidence in favor of such explanation. Section 6 quantifies the employment guarantee’s consequences on intra-household resource allocation. Section 7 discusses the empirical strategy and results in this paper as compared to those in related studies. Section 8 summarizes.

2. Data

Table 1 provides a self-contained summary of this section. I discuss essential aspects of the analysis samples and data available below. Appendix 1 provides more details.

2.1 Labor-Market Analysis

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. The 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).

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 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 use 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). 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. Panel a. of Appendix Table A.1 describes the samples based on the EU-NSS and IHDS. It

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)
Observations					
Level	Individual	Individual (each individual observed two times)	Husband-wife pairs	Husband-wife pairs (each pair observed two times)	Married women (each individual observed two times)
Household Role	Head, child of head, or child-in-law of head	Head, child of head, or child-in-law of head	Husband is the male head	Husband is the male head, child of head, or child-in-law of head	Female head, child of head, or child-in-law of head
Urban	259,351 men; 252,734 women	8,105 men; 7,445 women	131,463 pairs	6,866 pairs	4,347 women
Rural	415,902 men; 418,804 women	18,514 men; 16,829 women	221,159 pairs	16,299 pairs	8,940 women
Age Profile	25 to 64 years old at the time of the survey	25 to 64 years old at the time of both surveys	Husband 25 to 64 years old at the time of the survey	Husband 25 to 64 years old at the time of the survey	25 to 64 years old at the time of both surveys
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	Domestic independence index, body mass index, height
Sampling					
Data Set of Origin	EU-NSS	IHDS	HE-NSS	IHDS	IHDS
Sampling Design	Seven repeated cross-sections	Longitudinal	Seven repeated cross-sections	Longitudinal	Longitudinal
Periods of Observation	1999-2000, 2004, 2004-05, 2005-06, 2007-08, 2009-10, 2011-12	2004-2005 and 2011-2012	1999-2000, 2004, 2004-05, 2005-06, 2007-08, 2009-10, 2011-12	2004-2005 and 2011-2012	2004-2005 and 2011-2012
Representativeness	National for each period	National for the period 2004-2005	National for each period	National for the period 2004-2005	National for the period 2004-2005
States in the Sample	34 (balanced across periods)	21 (major states; balanced across periods)	34 (balanced across periods)	21 (major states; balanced across periods)	21 (major states; balanced across periods)
Districts in the Sample	582 (maximum across cross-sections; unbalanced across periods)	326 (balanced across periods)	582 (maximum across cross-sections; unbalanced across periods)	326 (balanced across periods)	326 (balanced across periods)

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. In all of the samples, I observe age, spouse age (if married), caste, religion, and marital status at the individual level. I merge year-specific, district-level agricultural (crop) volume and state-level (monsoon) rain volume available in Government of India (2022) into each individual or husband-wife observation.

indicates that the average individual demographics closely align across them.

In both samples, I observe age, caste, religion, socioeconomic disadvantage, and marital status.² I also merge in district-level agricultural and state-level rain information from Government of India (2022). The number of observations in the most basic specifications that I present below corresponds to the number of individuals in the households 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 not directly addressed in my empirical strategies.

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 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 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 who report 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 analyzed to 2018 purchasing-power-parity (PPP) dollars.

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-

²I 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”).

NSS that coincide in timing with the seven rounds of the EU-NSS described in Section 2.1. 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.2 indicates that the basic demographics of the sample based on the HE-NSS closely align with those described in Appendix Table A.1 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 also not a concern.

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. All consumption variables are monthly.³

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.

³The 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 nationally representative sample of 2004-2005, 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%.

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. I also observe body-mass index (BMI) and height.

3. The Mahatma Gandhi National Rural Employment Guarantee Act

The employment guarantee provides casual work in construction job sites. The work provided 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); workers are easily substitutable with one another. The employment guarantee provides up to 100 job days per household. Households are free to decide how these days are split between adult individuals. Ministry of Rural Development (2005a) states that the primary objective of the employment guarantee is to “enhance the livelihood security in rural areas.” Ministry of Rural Development (2005b) states secondary objectives, which include generating productive assets [infrastructure], protecting the environment, reducing rural-urban migration, and fostering social equity. Another of its secondary objective is empowering women by promoting their participation in the labor force.

District-Level Implementation Phases. Ministry of Rural Development (2005a) states that the large scale of the employment guarantee required a gradual implementation. 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 national social policies (Planning Commission, 2003). Prioritized districts were at a relative socioeconomic disadvantage by design. Their employment guarantee began in April 2006 (Phase 1) or April 2007 (Phase 2). In the rest of the districts, it began in April 2008 (Phase 3). The classification of districts by phase is available in Ministry of Rural Development (2010).⁴

⁴The 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).

Treatment and Control States. During the period that I analyze, most employment-guarantee jobs were provided in seven (treatment) states. I observe individuals from each of these treatment states in the samples based on the EU-NSS, IHDS, and HE-NSS. Imbert and Papp (2015) and Klonner and Oldiges (2022) explain that administrative capacity and experience in providing social programs determine the difference in job provision between the seven treatment states and the remaining states in India, referred to as control states henceforth.⁵ While these two studies use the treatment-control state classification to corroborate that their estimated impacts are driven by treatment states, I integrate this source of variation into my identification strategies. I thus 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).

Panel (a) of Figure 1 is based on the IHDS labor-market sample described in Section 2.1. I classify individuals into their district phase and state treatment status to compute the average annual employment-guarantee job days they took up in 2011-2012. The figure shows that it ranges between 11 and 15 days for women in treatment states, while it is about two for women in control states. The take-up in treatment states is relatively large, representing 11% to 16% of the average annual days worked across activities for women of 105 days observed before the employment guarantee (2004-2005). For men, take-up in 2011-2012 was relatively small compared to their average annual days worked across activities of 200 days before the employment guarantee or compared to the take-up of women.

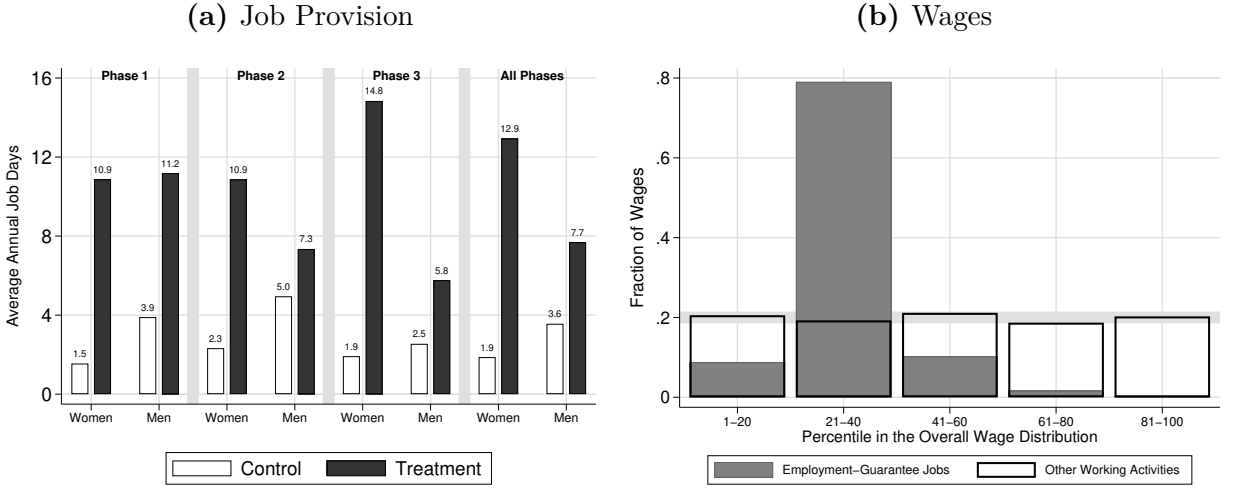
Panel (a) of Figure 1 could lead to the premature conclusion that the employment guarantee empowers women by providing them jobs at a higher rate than men. Analyzing the potential change in participation across all labor-market activities and the intra-household consequences of this potential change is necessary before such a conclusion.⁶

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. I do not exploit this variation in my empirical analysis

⁵In 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.

⁶Appendix Figure A.1 is analogous in format to Figure 1. However, it delimits the source sample to married individuals, who I focus on after Section 4. The figures are very similar.

Figure 1. Employment Guarantee: Provision and Wages



Note: Panel (a) displays the average annual days worked in employment-guarantee jobs (individuals who do no work in employment-guarantee jobs are assigned 0 days). The calculations are based on the 2011-2012 rural observations of the IHDS labor-market sample described in Section 2.1. It displays the average by district-level implementation phase and state treatment status. Panel (b) is based on the the 2011-2012 rural male observations of the IHDS labor-market sample described in Section 2.1. It plots the fraction of wages that fall into each of the quintiles of the overall distribution.

because the fixed effects and controls used below effectively incorporate it.

Panel (b) of Figure 1 uses the rural male 2011-2012 observations of the IHDS labor-market sample.⁷ I compare the wages in the employment guarantee and in any other activity with the overall wage distribution. The wages in any other activity uniformly fit into the quintiles of the overall distribution by design (most wages in India are in activities other than the employment guarantee). 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 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 (instead, their available employment-guarantee wages are also at the bottom of the distribution).

Corroborating the Employment-Guarantee Information in the IHDS. Four exercises corroborate that the rural subsample of the IHDS labor-market sample replicates aggregate moments from official records of the employment guarantee available in Government

⁷I 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 due to the high male labor participation rate (see Section 4.1).

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 31.⁸

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 fiscal-transparency reform reduces leakage. A similar reform to that in Bihar took place nationally in 2011.⁹ No evidence indicates that corruption compromises the individual or household average treatment effects discussed below. However, fund leaking 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

⁸The 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 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 2012-2012 and 2018-2019.

⁹The reform mandated the publication of live updates on funding and participation. The live updates appear in Government of India (2022).

bounds of a roll-out without corruption.

4. Labor Force Participation and the Employment Guarantee

4.1 Labor Force Participation Before and After the Employment Guarantee

The EU-NSS labor-market sample allows me to describe the aggregate context of labor force participation before and after the employment guarantee.¹⁰ This sample contains four periods before the start of the employment guarantee, which is convenient for the analysis of pre-policy trends. The sample also includes most of the decrease in female labor force participation observed during the last thirty years.¹¹

Panel (a) of Figure 2 displays the time series of the female labor force participation rate based on this sample. The rate remained virtually constant between 1999-2000 and 2005-2006. After that, it decreased by eight percentage points. This substantial decrease occurred while the female labor force participation rate was increasing worldwide, and the male labor force participation rate in India barely changed, which Panel (c) of Figure 2 illustrates. The decrease 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 (b) of Figure 2 shows that the overall decrease observed for women is driven by those who are rural and married. Among them, the decrease is more pronounced for the disadvantaged, who belong to the households *de facto* targeted by the employment guarantee. The decrease started right around its announcement. It is pertinent to examine if the employment guarantee contributes to the overall decrease.

4.2 Frameworks for Micro-Data Analysis


Event-Study Framework for Repeated Cross-Sectional Data. I partition the individuals in the EU-NSS labor-market sample into subsamples based on the implementation phase of their districts of residence, denoted by p , which establishes when the employment

¹⁰The EU-NSS labor-market sample is appropriate for this description because it has more frequent observations than its IHDS counterpart. It also has a longer time span. Two exercises further justify its use. First, Appendix Figures A.5 and A.6 show that temporal differences essential for the analysis below calculated in the EU-NSS sample are identical to their counterparts in the IHDS sample. Second, Appendix Figure A.2 shows that the female labor force participation time series obtained in it is virtually identical to the time series reported in World Bank (2022a).

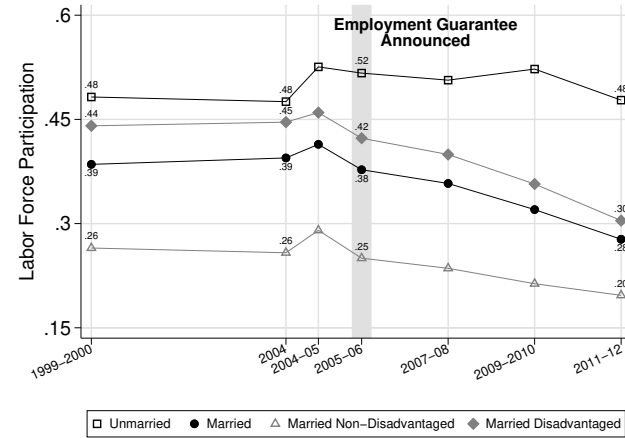
¹¹Appendix Figure A.2 shows that the additional decrease after 2012 is minor relative to the decrease observed in Figure 2. I do not use the three EU-NSS rounds before 1999-2000 because they do not contain geographic identifiers, an essential component of my analysis. In Appendix Figure A.3, I expand the sample 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.

Figure 2. Labor Force Participation in India

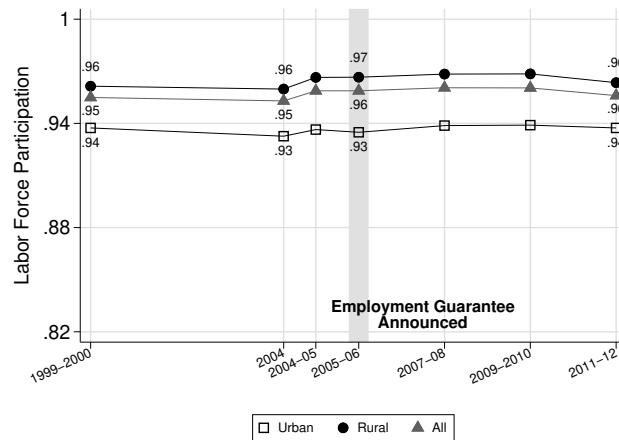
(a) Women by Sector




(b) Rural Women



(c) Men by Sector



(d) Rural Men



Note: Panel (a) displays the fraction of women who participated in the labor force during the year that the horizontal axis indicates. The calculation includes married and unmarried (never married, separated, divorced, or widowed) women who were between 25 and 64 years old during the corresponding year. Panel (b) 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 the rural married women into the participation rates of those disadvantaged and non-disadvantaged. Panels (c) and (d) are analogous in format to Panels (a) and (b) for men.

guarantee starts for them. Their districts belong to treatment or control states, which establishes the employment-guarantee intensity that they are exposed to (high intensity in treatment states; low intensity in control states). In each subsample, I estimate

$$\begin{aligned}
y_{ipg} &= \nu_i + \tau_g + \sum_{j < -1} \gamma_j^p \cdot \mathbf{1}[i \text{ lives in a treatment state}]_i \cdot \mathbf{1}[g = j]_g \\
&+ \gamma^p \cdot \mathbf{1}[i \text{ lives in a treatment state}]_i \cdot \mathbf{1}[g > 0]_g + \varepsilon_{ipg},
\end{aligned} \tag{1}$$

where y_{ipg} indicates if individual i , who resides in a Phase- p district, participates in the labor force during period g . g denotes event time (i.e., the number of periods after the start of the employment guarantee). ν_i is a generic fixed effect. It varies across the specifications that I estimate, though it always includes district fixed effects. τ_g is an event-time fixed effect. $\mathbf{1}[\cdot]_i = 1$ if the statement in brackets is true and $\mathbf{1}[\cdot]_i = 0$ otherwise. The same definition applies to $\mathbf{1}[\cdot]_g$. ε_{ipg} is an error term.

Difference-in-Difference Estimands. γ_g^p and γ^p are the period-specific (conditional) labor force participation rates 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 because the fully saturated version of Equation (1) is not identified. This specification implies that the coefficients γ_g^p and γ^p are relative to the treatment-control difference in the reference period (i.e., they are difference in differences).

Quasi-experimental Variation. I combine district (timing) and state (intensity) sources of quasi-experimental variation in this framework. The timing index g depends on an individual's district of residence. For individuals residing in Phase-1 or Phase-2 districts, I observe four periods before the employment guarantee starts in addition to a reference period. I also observe two periods after it starts. The before, reference, and after periods occur in the same calendar years for those in Phase-1 or Phase-2 districts. Therefore, I estimate Equation (1) pooling districts in these two phases, which amounts to assuming that $\gamma_g^1 = \gamma_g^2$ for every g and $\gamma^1 = \gamma^2$. In this case, I label g with the midpoint between the two phases. For Phase-3 districts, I observe five periods before the employment guarantee, a reference period, and one period after.¹²

Parameter of Interest and Identification Assumptions. The parameter of interest is

¹²In the sample of individuals residing in Phase-1 or Phase-2 districts, I group the two periods after the employment guarantee starts and label the binned period according to the binned midpoint of the two phases. In Phase-1 or Phase-2 districts, g takes values from the set $\{-8.5, -4.5, -3.5, -2.5, -0.5, 2.5\}$. In the sample of individuals residing in Phase-3 districts, it takes values from the set $\{-10, -6, -5, -4, -2, 0, 2\}$. The reference periods are -0.5 and 0 . Appendix 4 explains and tabulates the mapping between the calendar years observed in the sample and the index g .

the coefficient associated with the period after the employment guarantee starts: γ^p . It is a treatment-control comparison, or, more precisely, a “high-dose to low-dose of treatment” comparison. This comparison is an absolute-term lower bound of the ideal “high-dose to no-dose of treatment” comparison (Heckman et al., 2000). This is an important caveat of the estimates presented throughout the paper: I can only identify and estimate the referred lower bounds.

My empirical design is akin to that in Sun and Abraham (2021). γ^p is an estimator of the average treatment on the treated (ATT) in districts that began implementation in Phase p . It 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 $g < -1$. If this implication holds, parallel trends before implementation hold, which favors parallel trends after implementation.

The ATT in this paper identifies the effect of the employment guarantee’s intention to treat. I aim to understand the impact of the sole existence of the insurance provided by the employment guarantee. Other parameters more directly focus the actual take-up of employment-guarantee jobs. These parameters are not the focus of this study.

An Aggregate ATT. I aggregate γ^p across phases using the estimator

$$\text{WDiD} := \sum_{p=1}^3 [\text{fraction of individuals in Phase-}p \text{ districts}] \cdot \gamma^p, \quad (2)$$

which is a population-weighted sum of the phase-wise difference-in-difference ATT estimators. In practice, estimating the WDiD requires two γ^p coefficient estimates, given the assumption that $\gamma^1 = \gamma^2$, and the population weights, which I estimate using their sample counterparts.

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:

$$y_{ig} = \nu_i + \tau_{\text{after}} + \gamma_{\text{after}} \cdot \mathbf{1}[i \text{ lives in a treatment state}]_i \cdot \mathbf{1}[g = \text{after}]_g + \varepsilon_{ig} \quad (3)$$

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 in ν_i individual fixed effects (which subsume district fixed effects). 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.

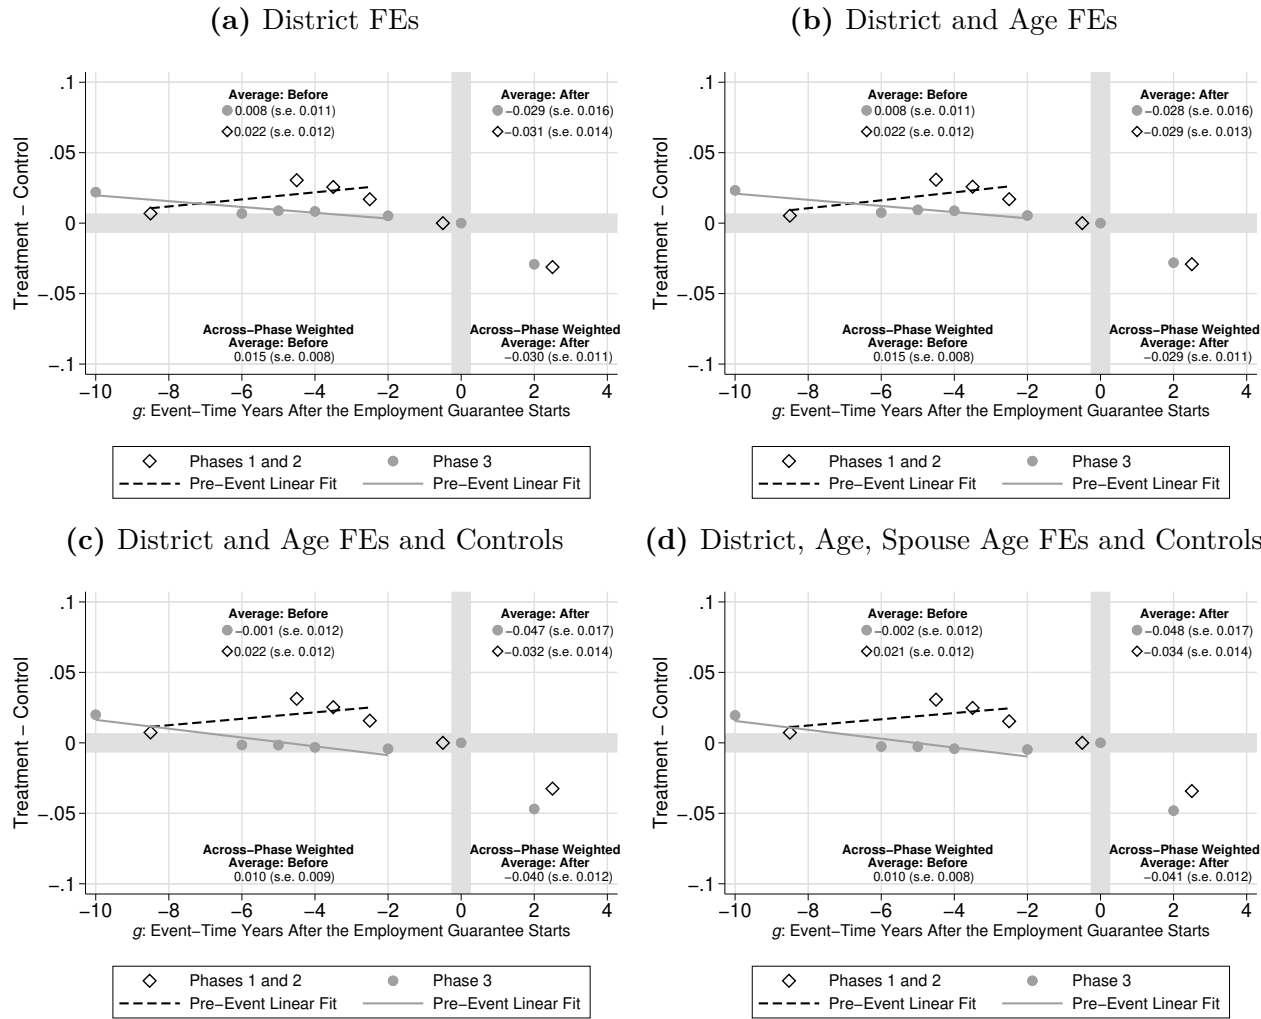
4.3 Estimates

Rural Married Women. I first analyze the subsample of rural married women, which are the great majority of women in the rural Indian population. I estimate the coefficients γ_g^p and γ^p in Equation (1) for four different specifications of ν_i and display them in Figure 3. The simplest specification includes only district fixed effects, the minimal requirement for identification given the quasi-experimental variation used. The other specifications progressively add age fixed effects, linear agricultural and rain controls described in Section 2, and spouse age fixed effects. Across specifications, there is evidence in favor of $\gamma_g^p = 0$ for $g < -1$, an implication of the assumptions by which γ^p identifies the ATT. The estimates of the coefficients γ_g^p are close to 0; their average is small in magnitude and in no case statistically differs from 0 when using standard significance levels. The figure also shows that estimates of the phase-wise ATTs are relatively homogenous. Their population-weighted average, representing the aggregate ATT, indicates that the employment guarantee reduces female labor force participation between 2.9 (s.e. 1.1) and 4.1 (s.e. 1.2) percentage points.

Panel a. of Table 2 displays the details of the aggregate ATT estimates arising from Figure 3. Panel b. of Table 2 displays their counterpart estimates based on the DiD estimator of Equation (3). I estimate specifications analogous to those considered when estimating Equation (1), except that individual fixed effects replace district fixed. Estimates are very similar across samples, reinforcing that within-individual policy responses drive them.

Basic Identification Threats. An important limitation of the evidence based on the EU-NSS labor-market sample is that it relies on within-district variation for identifying a parameter describing individual behavior. The estimate based on the IHDS labor-market sample ameliorates this issue, as it relies on within-individual variation. However, the short time span of this sample does not allow me to justify the identification assumptions in it. Therefore, the justification and visual evidence are based on the EU-NSS labor-market

Figure 3. Labor Force Participation of Rural Married Women and the Employment Guarantee, Event Studies



Note: Panel (a) displays event-study coefficient estimates based on Equation (1). The estimates are the difference in the (conditional) labor force participation rate between the districts in treatment and control states. Districts either belong to implementation Phases 1 and 2 or Phase 3. All treatment and control states have districts in the three implementation phases. The rates are conditional on district fixed effects. The average and corresponding standard error are in the display for the pre-event rate differences. For the post-event period, the average differences are estimates of the average treatment on the treated. In both cases, the population-weighted average and corresponding standard error across phases are in the display. The post-event population-weighted average is an estimate of the aggregate average treatment on the treated in Equation (2). The standard errors (s.e.) are bootstrapped clustering at the state \times age-group level. Panels (b) to (d) are analogous in format to Panel (a), conditioning the participation rates on the fixed effects and controls indicated in the label. **Sample:** Rural married female subsample of the EU-NSS labor-market sample.

Table 2. 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	-0.030	-0.029	-0.040	-0.041	0.000	0.002	0.001	0.001
(s.e.)	(0.011)	(0.011)	(0.012)	(0.012)	(0.004)	(0.003)	(0.004)	(0.004)
Observations	366,793	366,793	331,586	331,486	367,317	367,317	331,540	331,494
<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
(s.e.)	(0.025)	(0.020)	(0.020)	(0.020)	(0.006)	(0.006)	(0.006)	(0.006)
Observations	13,375	13,375	13,154	13,140	16,328	16,328	16,078	16,048

Note: Column (1) of Panel a. displays details from the estimation of the aggregate average treatment on the treated based on Equation (2) for rural married women. The required estimates of the average treatment on the treated for each phase are based on Equation (1). ν_i includes district (Dist) fixed effects when estimating this equation. Columns (2) to (4) are analogous in format to Column (1). Their only difference is the inclusion of different fixed effects or controls in the specification of ν_i , which is indicated in the column labels. 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 (3). Panel b. is based on longitudinal data rather than repeated cross-sections. It thus replaces district with individual (Indv) fixed effects in the specification of ν_i . The standard errors (s.e.) are bootstrapped clustering at the state \times age-group level. **Sample:** Rural married female (left) and male (right) subsamples of the EU-NSS (Panel a.) and IHDS (Panel b.) labor-market samples.

sample.

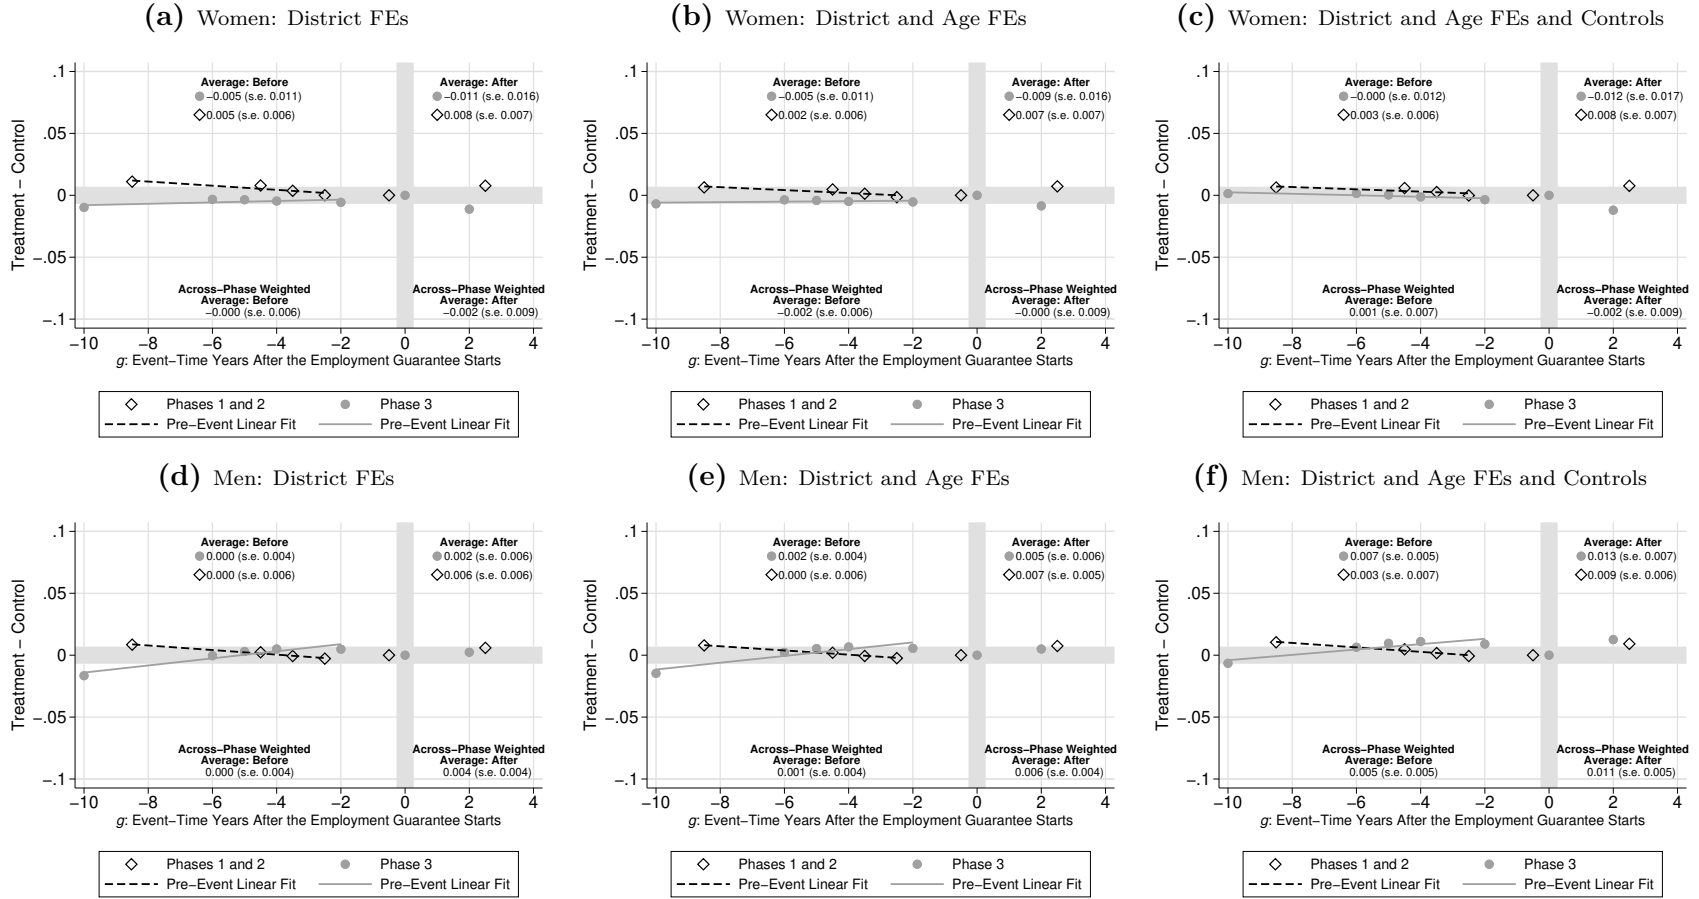
In either framework, identifying the parameter of interest requires a time-invariant classification of the district and state where individuals reside. Rural-rural, across-district migration would compromise my identification strategy and could be a symptom of anticipative behavior. For example, individuals 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, and, therefore, invalidate the interpretation of the negative impact I provide below. 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. Other issues related to the incidental truncation of unmarried women could also bias the estimates in Table 2. Figure 4 is analogous in format to Figure 3. The outcome is “being married,” as opposed to single (never married), divorced, or widowed. I estimate specifications for the women and men in the EU-NSS labor-market sample.¹³ The resulting estimates of the aggregate ATT are precisely estimated at 0, except for the isolated case of Panel (f). I thus interpret the results as supporting a lack of impact on being married.

Another threat could be the existence of pre-event trends. Figure 3 shows a positive trend for Phases 1 and 2 driven by periods -8.5 and -3.5 . However, the trend reverts to essentially 0 after period -2.5 . For Phase 3, the figure shows a negative trend, though all pre-trend point estimates are essentially 0 except for a relatively small, positive point estimate at -10 . I argue that these trends are not concerning for the following reasons. First, I cannot reject the null hypothesis that the average pre-trend point estimate is 0, as noted before. Second, the trends have opposite directions across phases. They do not

¹³Figure 4 does not include specifications with spouse fixed effects because the sample includes all women and men, not only those who are married. Appendix Table A.7 is analogous in format to Table 2, including estimates of the aggregate ATT based on the EU-NSS and the IHDS labor-market samples. The estimates closely align across samples.

Figure 4. Marital Status of Rural Women and Men and the Employment Guarantee, Event Studies



Note: Panel (a) displays event-study coefficient estimates based on Equation (1) with “being married” as an outcome, as opposed to single (never married), divorced, or widowed. The estimates are the difference in the (conditional) marriage rate between the districts in treatment and control states. Districts either belong to implementation Phases 1 and 2 or Phase 3. All treatment and control states have districts in the three implementation phases. The rates are conditional on district fixed effects. The average and corresponding standard error are in the display for the pre-event rate differences. For the post-event period, the average differences are estimates of the average treatment on the treated. In both cases, the population-weighted average and corresponding standard error across phases are in the display. The post-event population-weighted average is an estimate of the aggregate average treatment on the treated in Equation (2). The standard errors (s.e.) are bootstrapped clustering at the state \times age-group level. Panels (b) to (c) are analogous in format to Panel (a), conditioning the participation rates on the fixed effects and controls indicated in the label. Panels (d) to (f) are analogous in format to Panels (a) to (c) for the sample of rural men. **Sample:** Rural female and male subsamples of the EU-NSS labor-market sample.

represent a general downward trend aligning with their corresponding average treatment on the treated estimate. Third, Appendix Figure A.7, analogous in format to Figure 3, shows linear predictions of the average treatment on the treated for Phases 1 and 2 and Phase 3 based on the pre-event trends. I reject the null hypothesis that their corresponding predicted aggregate ATT is equal to the actual point estimate of the aggregate ATT in Panel (a) of Figure 3. Appendix Figure A.7 shows that, for the other three specifications, I also reject this null hypothesis.¹⁴

Rural Married Men. Appendix Figure A.9 is analogous in format to Figure 3. It confirms the evidence in Table 2 indicating that the employment guarantee has no impact on the labor force participation for rural married men. Section 5 explains the lack of impact in conjunction with the negative impact on the participation of rural married women.

Rural Unmarried Women and Men. The event-study evidence in Appendix Figure A.10 and corresponding ATT estimates in Appendix Table A.6, which are based on the EU-NSS labor-market sample, indicate that the employment guarantee had no impact on unmarried women and men. The only point estimate that differs from 0 in magnitude has a large standard error—0.015 (s.e. 0.011). The other estimates range between -0.002 (s.e. 0.017) and 0.007 (s.e. 0.010). In the IHDS labor-market sample, only 1,788 women and 850 men remain unmarried during the two observation periods. Thus, the estimates are less reliable. They indicate a positive ATT for women and a negative ATT for men. However, the specificity of the sample makes these latter results less conclusive. The conclusions regarding unmarried individuals are not definitive. Further investigation of this minority is outside the scope of this paper. I focus on married individuals henceforth.

Falsification Tests. I consider two placebo subsamples for providing falsification tests. The first includes rural, married, 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 are 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

¹⁴Appendix Figure A.8 presents an additional robustness check. It is analogous in format to Figure 3 except that it does not consider observations from 2004-2005. The motivation of this exercise is the labor force participation blip observed for this period in Figure 2. The objective is to discard that such a blip drives the results from the event-study framework. In this exercise, the pre-event point estimates or their averages barely change. The corresponding estimates of the average treatment effect remain very similar because both the reference and post-event periods do not change when dropping the 2004-2005 observations.

areas.¹⁵ By construction, the placebo subsamples and the subsample of rural married women differ in observed and unobserved characteristics. However, the placebo subsamples provide an additional assessment of the treatment and control trends in a no-implementation scenario *after* the employment guarantee starts. Table 3 displays the impacts for both placebo subsamples, based on the EU-NSS and the IHDS labor-market samples. It shows no impact in both cases.¹⁶

Inference. The standard errors presented throughout the paper are the standard deviation of the empirical bootstrap distribution of the corresponding estimates. They account for sampling variation in all estimation stages—e.g., they account for sampling variation in both the phase-wise average treatment on the treated and the weights in Equation (2). The sampling is clustered at the state \times age-group level when obtaining the empirical bootstrap distributions. Appendix 10 justifies this clustering.¹⁷ It documents that, in this context, asymptotic rules of thumb for significance apply (e.g., if the estimate divided by its standard error is greater or equal to 1.96, the null hypothesis that the corresponding impact is 0 is rejected with a significance level of 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 in this section (i.e., the negative impact on the labor force participation of rural married women). The calculation requires the following components. First, the share of the rural population to which the negative impact applies—in this case, the number of women who are rural, married, and residing in treatment states as a fraction of all rural married

¹⁵I am able to implement the event-study framework in the urban subsample because districts have urban and rural areas.

¹⁶The results in Table 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). Appendix Figure A.11 displays the event-study evidence corresponding to the placebo subsamples in the EU-NSS labor-market sample. The pre-trends are imprecise, which is likely due to small sample size. However, the ATT estimates are precise at indicating a lack of impact.

¹⁷I observe individuals of forty ages ($65 - 24 + 1$). In the EU-NSS, I form 10 equally-sized age groups (each group contains four ages). The corresponding state \times age group bins are the clusters. In the IHDS, which includes a subsample of the states in India, I follow an analogous procedure forming 15 equally-sized age groups. I thus observe a similar number of clusters across the samples based on the two data sets and consistently obtain similar standard errors in them. The clustering is a mid-point between a common state \times age clustering in quasi-experimental designs that impact women of different ages within clustered states (e.g., Goldin and Katz, 2002) and the more conservative approach of clustering at the state level. The former approach parameterizes variance differences among women differently affected by policies related to labor force participation and similar outcomes given their stage in the life cycle. Appendix 10 documents that the statistical hypotheses tested throughout the paper jointly favor the economic significance of the mechanisms driving them. Appendix 10 also discusses inference based on non-parametric, score, and wild-bootstrap p -values, as well as recent analyses of standard-error clustering (e.g., Abadie et al., 2022).

Table 3. Labor Force Participation of Women in Placebo Samples and the Employment Guarantee, Estimates of the Average Treatment on the Treated

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
<i>Panel a. Data: NSS-EU; Year Span: 1999-2000 to 2011-2012; Estimator: WDiD</i>								
	Urban				Rural Non-Disadvantaged			
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	-0.006	-0.006	-0.008	-0.007	-0.006	-0.006	0.001	0.000
(s.e.)	(0.010)	(0.010)	(0.010)	(0.010)	(0.014)	(0.014)	(0.015)	(0.015)
Observations	213,660	213,660	186,212	186,173	108,115	108,115	103,287	103,252
<i>Panel b. Data: IHDS; Year Span: 2004-2005 and 2011-2012; Estimator: DiD</i>								
	Urban				Rural Non-Disadvantaged			
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	0.026	0.026	0.010	0.010	-0.015	-0.015	-0.020	-0.024
(s.e.)	(0.015)	(0.015)	(0.016)	(0.017)	(0.022)	(0.023)	(0.024)	(0.024)
Observations	6,257	6,257	6,064	6,055	3,587	3,587	3,544	3,537

Note: Column (1) of Panel a. displays details from the estimation of the aggregate average treatment on the treated based on Equation (2) for urban married women. The required estimates of the average treatment on the treated for each phase are based on Equation (1). ν_i includes district (Dist) fixed effects when estimating this equation. Columns (2) to (4) are analogous in format to Column (1). Their only difference is the inclusion of different fixed effects or controls in the specification of ν_i , which is indicated in the column labels. Columns (5) to (8) are analogous in format to Columns (1) to (4) for rural non-disadvantaged married women. Panel b. is analogous in format to Panel a. The estimate of the aggregate average treatment on the treated is based on Equation (3). Panel b. is based on longitudinal data rather than repeated cross-sections. It thus replaces district with individual (Indv) fixed effects in the specification of ν_i . The standard errors (s.e.) are bootstrapped clustering at the state \times age-group level. **Sample:** Urban married (left) and rural married non-disadvantaged (right) female subsamples of the EU-NSS (Panel a.) and IHDS (Panel b.) labor-market samples.

women. Second, an estimate of the negative impact (i.e., an estimate of the ATT, noting that the ATT is an annual impact for a given year after the employment guarantee). Third, the aggregate decrease used to contrast the magnitude of the ATT. I provide two calculations: one in which the aggregate decrease is that observed between 2005-2006 and 2007-2008 (i.e., the immediate decrease after implementation), and the other in which the aggregate decrease is the average difference between the rates observed for 2005-2006 and before and the rates observed for 2007-2008 and after. An example of the first calculation is:

$$\left(\underbrace{0.22}_{\text{share of rural married treated among all rural women}} \cdot \underbrace{0.041}_{\text{average treatment on the treated}} \right) / \underbrace{0.031}_{\text{aggregate decrease between 2005-2006 and 2007-2008}} \approx \underbrace{0.30}_{\text{decrease accounted for by the employment guarantee}} \quad (4)$$

Applying this calculation using the estimates of the ATT in Table 2 and the two possible aggregate decreases, the employment guarantee accounts for 10% to 30% of the decrease observed for rural women in Panel (a) for Figure 2.

This calculation is “back-of-the-envelope” for two reasons. First, it is a prediction about a level change rather than an estimate of a slope like the ATT. It thus assumes that the ATT applies homogeneously in treatment states and, as a level, is a fraction of the decreasing trend. Second, it is valid for an arbitrary post-implementation period rather than the entire period in which the policy is in place. Despite these caveats, the calculation provides aggregate context, indicating that a policy treating only about 22% of the population accounts for up to 30% of a countrywide decrease.

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

Work by Activity and the Intensive and Extensive Margins of Labor Force Participation. The labor-market sample of the IHDS, along with sources cited throughout the paper, indicate a substantial take-up of employment-guarantee jobs. This take-up is salient for rural women, most of whom are married. The negative impact of the employment guarantee on their labor force participation is thus puzzling. I clarify this apparent contradiction by estimating the impact of the employment guarantee on days worked in mutually exclusive and exhaustive categories of work. I use the DiD estimator in Equation (3) and longitudinal data on days worked by activity available in the IHDS labor-market sample.¹⁸ Panel (a)

¹⁸I estimate Equation (3) for each relevant dependent variable using the most complete specification of ν_i , which includes individual, age, and spouse age fixed effects as well as controls. Days worked in employment-guarantee jobs are set to 0 for all of the observations in 2004-2005 (before the employment guarantee). Appendix Figures A.12 to A.14 are analogous in format to Figure 5. They use the other three specifications

of Figure 5 shows the estimates. The average treatment-control difference in the annual take-up of employment-guarantee jobs is 10.3 days. However, the employment guarantee also decreases women’s agricultural and non-agricultural work by annual averages of 12.2 and 13.4 days, from baseline averages of 40.8 and 20.2. These negative impacts outweigh the average treatment-control difference in employment-guarantee take-up; they generate the negative average impact of 14.8 on annual days worked across activities from a baseline average of 101.8.¹⁹

The same estimator and data source allow me to clarify that the policy affects labor force participation of rural married women at the intensive and extensive margins. I estimate the impact of the employment guarantee on mutually exclusive and exhaustive categories of annual days worked. Panel (c) of Figure 5 displays the results. The employment guarantee increases the probability of working 0 days per year (i.e., quitting the labor force) by 0.039, a result which is summarized in Section 4. Further, it increases the probability of working 0 to 90 days per year by 0.09 by shifting the distribution of days worked to the left. The shift occurs through reductions of 0.05 and 0.04 in the probabilities of working 91 to 180 days per year and 181 to 365 days per year. About $(0.039/0.09) \cdot 100 \approx 43\%$ of women whose annual days worked shift to the left due to the employment guarantee quit the labor force. This high quitting rate is expected given the context-specific gender role of women as added workers.

Panels (b) and (d) of Figure 5 are analogous in format to Panels (a) and (c) for rural married men (husbands). 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 of annual days worked. Mechanically, reallocation across activities explains the simultaneous reduction in days worked and the massive aggregate provision of employment-guarantee jobs. An economic explanation follows.

5.1 The Employment Guarantee as Insurance of Household Earnings

I assume that 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 .²⁰ The


of ν_i considered throughout the paper and show very similar results.

¹⁹The baseline average is the 2004-2005 (before the employment guarantee) average of annual days worked in treatment states. For rural married men, the baseline averages of days worked in self-employment, agricultural-wage jobs, and non-agricultural-wage jobs are 77.3, 57.3, and 83.9.

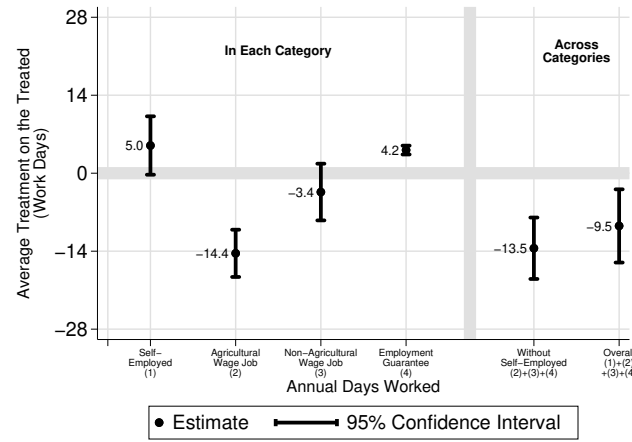
²⁰This assumption holds if the utility derived from such activities is greater than their cost. For example, a woman may derive utility from raising her children (e.g., she enjoys raising her children because it allows

Figure 5. Annual Days Worked and the Employment Guarantee

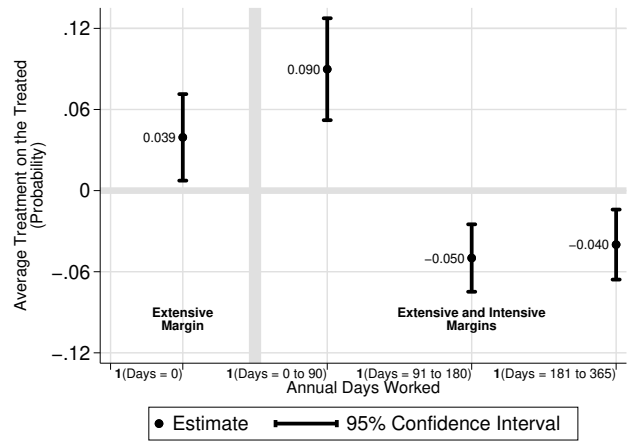
(a) Annual Days Worked By Activity, Women




(b) Annual Days Worked By Activity, Men



(c) Overall Days Across Activities, Women



(d) Overall Days Across Activities, Men



Note: Panels (a) and (c) display estimates of the aggregate average treatment on the treated based on Equation (3) for each of the dependent variables labeled in the horizontal axes. These panels are based on the subsample of rural married women. Days worked are measured annually. Individuals who do no work a certain category are assigned 0 days (i.e., days worked are not conditional on participation). The specification of ν_i includes individual, age, and spouse age fixed effects as well as controls. The standard errors (s.e.) and confidence intervals are bootstrapped clustering at the state \times age-group level. Panels (b) and (d) are analogous in format to Panels (a) and (c) for the subsample of rural married men.
Sample: Rural married female (a and c) and male (b and d) subsamples of the IHDS labor-market sample.

woman and her husband have $365 \times 2 = 730$ days per year, which they allocate between H and their total annual days worked, D . Phenomena such as involuntary unemployment due to draughts may restrict their work days. If risk averse, they accumulate a buffer stock (precautionary savings) in anticipation of such phenomena. Without phenomena restricting their work days, they choose the bundle (C^*, H^*) . In the presence of such phenomena, they work more than $D^* := 730 - H^*$ to accumulate a buffer stock. If their utility function is concave in C and H , the buffer stock is not financed solely by increasing D . Instead, the household reduces both C and H to make the reduction from (C^*, H^*) less extreme. For simplicity, I abstract from fully specified dynamic considerations and describe choices for a given period. I then explore a change in these choices upon the (permanent) implementation of the employment guarantee.

In the absence of the employment guarantee, restrictions do not occur every year. The household increases D in anticipation of future (stochastic) restrictions. This increase includes the primary worker as well as the secondary worker, who may not work at all in the absence of anticipated restrictions. Once the employment guarantee is in place, the risk of not reaching D^* is eliminated,²¹ shutting down a reason for accumulating the buffer stock. Figure 5 provides evidence consistent with this framework. It indicates that, once the policy is in place, there is some take-up of employment-guarantee jobs by those requiring them to reach D^* . However, all other (non-shocked) households drive an overall average decrease in days worked, from the level allowing them to accumulate the buffer stock in the absence of the policy towards D^* . When using the estimates in Figure 5, the observed average decrease amounts to 24.2 (total of average reductions for women and their husbands). Appendix Figure A.15 displays inference for this average, which has a standard error of 6.8.

I also test the implication of the framework that total household consumption increases.²² Figure 6 is analogous in format to Figure 3. It displays event-study evidence using the subsample of rural households of the HE-NSS consumption sample. It shows a positive impact on the log of household consumption per capita that differs statistically from 0 when using standard significance levels. Table 4 provides corroboration for this impact. It is analogous in format to Table 2. It shows that the estimates of the ATT based on the WDiD estimator of the event-study framework are essentially identical to the esti-

her to satisfy social norms or because she finds fulfillment in doing so) but also find it costly. If the net gain from this activity is strictly concave, the framework's assumption regarding non-market activities is valid.

²¹If the days of work provided by the employment guarantee are not sufficient to reach this level, the woman and her husband would still decrease their days worked and increase their consumption, getting closer to (C^*, H^*) relative to the scenario without the policy.

²²An additional implication of this framework is that household earnings from labor income decrease. This implication has empirical support: 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).

mates of the ATT based on the DiD estimator in Equation (3) and longitudinal data on log household consumption per capita available in the consumption sample of the IHDS. The estimates indicate that the employment guarantee increases monthly household consumption per capita by an average of 6% to 8% from a baseline average of 256 (2018 USD, PPP). If household consumption is used as a metric of well-being, the employment guarantee reduces household-level absolute poverty. A within-household distributional analysis follows.

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.²³ 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: $g \in \{\text{before, after}\}$ and treatment status: $d \in \{\text{control, treatment}\}$. These two indices define four regimes. The household allocates total consumption of goods, C_g^d , by solving

$$\max_{\mathbf{z}_g^d} \tilde{U}_g^d [U_g^{d,w}(\mathbf{x}_g^{d,w}), U_g^{d,h}(\mathbf{x}_g^{d,h})] \quad (5)$$

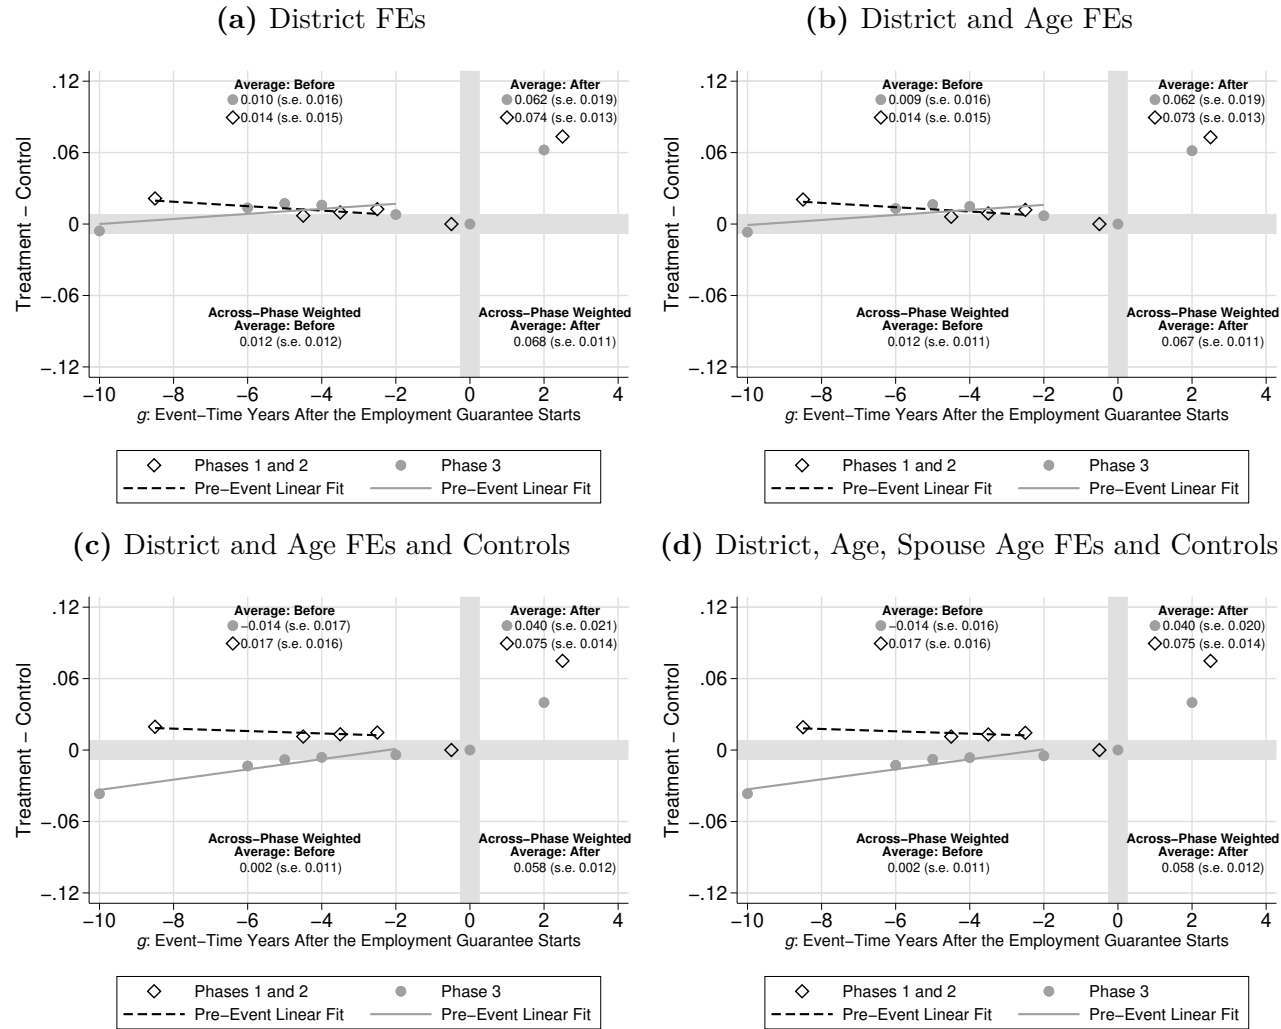
subject to

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

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 $r \in \{w, h\}$.

²³As before, I abstract from dynamic considerations and compare within-household allocations between a given period without the employment guarantee and a given period after the (permanent) policy change.

Figure 6. Log of Household Consumption per Capita and the Employment Guarantee, Event Studies



Note: Panel (a) displays event-study coefficient estimates based on Equation (1). The estimates are the difference in the (conditional) log of household consumption per capita between the districts in treatment and control states. Districts either belong to implementation Phases 1 and 2 or Phase 3. All treatment and control states have districts in the three implementation phases. The log of household consumption per capita is conditional on district fixed effects. The average and corresponding standard error are in the display for the pre-event average differences. For the post-event period, the average differences are estimates of the average treatment on the treated. In both cases, the population-weighted average and corresponding standard error across phases are in the display. The post-event population-weighted average is an estimate of the aggregate average treatment on the treated in Equation (2). The standard errors (s.e.) are bootstrapped clustering at the state \times age-group level. Panels (b) to (d) are analogous in format to Panel (a), conditioning log household consumption per capita on the fixed effects and controls indicated in the label. **Sample:** Subsample of rural households of the HE-NSS consumption sample.

Table 4. Log of Household Consumption per Capita of Rural Households and the Employment Guarantee, Estimates of the Average Treatment on the Treated

	(1)	(2)	(3)	(4)
<i>Panel a. Data: HE-NSS; Year Span: 1999-2000 to 2011-2012; Estimator: WDiD</i>				
Fixed Effects	Dist	Dist, Age	Dist, Age	Dist, Age, Spouse Age
Controls	No	No	Yes	Yes
Estimate	0.068	0.067	0.058	0.058
(s.e.)	(0.011)	(0.011)	(0.012)	(0.012)
Observations	221,159	221,159	198,308	198,294
<i>Panel b. Data: IHDS; Year Span: 2004-2005 and 2011-2012; Estimator: DiD</i>				
Fixed Effects	Indv	Indv, Age	Indv, Age	Indv, Age, Spouse Age
Controls	No	No	Yes	Yes
Estimate	0.073	0.078	0.056	0.055
(s.e.)	(0.017)	(0.015)	(0.016)	(0.016)
Observations	16,285	16,285	16,036	16,034

Note: Column (1) of Panel a. displays details from the estimation of the aggregate average treatment on the treated based on Equation (2). The required estimates of the average treatment on the treated for each phase are based on Equation (1). ν_i includes district (Dist) fixed effects when estimating this equation. Columns (2) to (4) are analogous in format to Column (1). Their only difference is the inclusion of different fixed effects or controls in the specification of ν_i , which is indicated in the column labels. Panel b. is analogous in format to Panel a. The estimate of the aggregate average treatment on the treated is based on Equation (3). Panel b. is based on longitudinal data rather than repeated cross-sections. It thus replaces district with male household head (Indv) fixed effects in the specification of ν_i . The standard errors (s.e.) are bootstrapped clustering at the state \times age-group level. **Sample:** Rural subsamples of the HE-NSS (Panel a.) and IHDS (Panel b.) consumption samples.

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.²⁴ 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, 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

$$\Omega_g^{d,r}(C_g^d) = \eta_g^{d,r}(C_g^d) \cdot \omega_g^{d,r}(\eta_g^{d,r}(C_g^d) \cdot C_g^d) \quad (6)$$

²⁴Suppose 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. 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 a given price vector \mathbf{p}_g^d .²⁵ $\Omega_g^{d,r}(C_g^d)$ 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).²⁶

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 certain 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 (6) as:

$$\Omega_g^{d,r}(C_g^d) = \underbrace{\alpha_g^{d,r}}_{\text{constant}} + \underbrace{b_g^{d,r}}_{\text{slope}} \cdot \log(C_g^d) + \xi_g^{d,r}, \quad (7)$$

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 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. The larger the woman’s response relative to the overall household response, the

²⁵I suppress prices because I do not rely on them for identification. This is an advantage, as identification of resource shares in collective models usually requires price variation (Chiappori and Mazzocco, 2017).

²⁶For 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 $0.5 \cdot 0.10 = 0.05$.

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 (7) across w and h :

$$\begin{aligned} \Omega_g^d(C_g^d) &=: \Omega_g^{d,w}(C_g^d) + \Omega_g^{d,h}(C_g^d) \\ &=: \underbrace{\alpha_g^d}_{\text{constant}} + \underbrace{\beta_g^d}_{\text{slope}} \cdot \log(C_g^d) + \underbrace{\xi_g^d}_{:=\xi_g^{d,w} + \xi_g^{d,h}}. \end{aligned} \quad (8)$$

Estimation. Ordinary least-square estimation of Equation (7) for $r \in \{w, h\}$ and Equation (8) 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 $\log(C_g^d)$.²⁷ 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 $\left[\left(\eta_{\text{after}}^{\text{treatment},w} - \eta_{\text{after}}^{\text{control},w} \right) - \left(\eta_{\text{before}}^{\text{treatment},w} - \eta_{\text{before}}^{\text{control},w} \right) \right]$.

Structural Parameters and the Employment-Guarantee Impact On Them. Panels a. and b. of Table 5 summarize the assignable private good composites for women and husbands. They display the annual 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.²⁸ 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

²⁷The 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, which I annualize by multiplying by 12. The alternative measure used as an instrument is based on a variable directly measuring overall annual total household consumption. The instrumental-variable strategy tackles measurement error. It does not tackle more general concerns related to endogeneity.

²⁸The difference-in-difference implied by the averages in Table 5 is 2.1 (s.e. 0.5) for women and 5.7 (s.e. 1.3). The difference of 5.7 - 2.1 = 3.6 (s.e. 1.2) differs statistically from 0 when using standard significance levels.

guarantee in such a relative slope drives the ATT on the parameter.

Table 5. Summary 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
24.079	22.422	26.774	27.250
Panel b. Average Husband Private Assignable Composite (2018 USD, PPP)			
Before, Control	Before, Treatment	After, Control	After, Treatment
38.747	43.570	41.376	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 (0.005)	0.450 (0.007)	-0.039 (0.008)	85,161
Panel d. Identification (Rank) Test: Engel-Curve Slope			
Before, Control	Before, Treatment	After, Control	After, Treatment
0.023 (0.002)	0.022 (0.003)	0.019 (0.002)	0.018 (0.003)
Panel e. Identification (Rank) Test: First Stage F-stat			
Before, Control	Before, Treatment	After, Control	After, Treatment
48,900.2	34,532.9	13,417.4	10,294.4

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 employment guarantee. The units are annual 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 the employment guarantee (2004-2005), as well as an estimate of the aggregate average treatment on the treated (ATT) 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 (8) for treatment and control states before (2004-2005) and after (2011-2012) the employment guarantee. 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 (7) for $r \in \{w, h\}$ and Equation (8). In Panels c. and d., the standard errors (in parentheses) are based on the bootstrap clustering at the state \times age-group level. The F statistics are asymptotic and clustered at the state \times age-group level. **Sample:** Subsample of rural households of the HE-NSS consumption sample, limited to the observations in rounds 61 (2004-2005, before the employment guarantee) and 68 (2011-2012, after the employment guarantee).

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. 0.008). That is, it decreases the female share by 9% from the treatment-state baseline. This impact more than doubles the gap between

the intra-household share of resources of the woman and that of her husband, which was observed before the employment guarantee. It implies that, within the household, female bargaining power decreases as a result of the employment guarantee.

Panel d. of Table 5 provides estimates of β_g^d and inference on them. For the four regimes, I reject the null hypothesis $\beta_g^d = 0$, which is necessary for identification of $\eta_g^{d,w}$. Panel e. provides the standard first-stage instrumental-variable rank test. For each regime, I report the F statistic from comparing two models. The first model regresses $\log(C_g^d)$ on a constant and the alternative measure of total household consumption. The second model only includes a constant. The rule of thumb is that the rank condition is satisfied if the F statistic is larger than 10. The results indicate that the rank condition holds. The F statistics are very large, which is expected given that the two measures of total household consumption closely track each other (note that, in my instrumental-variable strategy, estimation of Equation (7) for $r \in \{w, h\}$ and Equation (8) has the same first stage). The two identification tests for each regime support that the model is identified across regimes.

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 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.²⁹ I test whether the employment guarantee generates this deterioration using the DiD estimator in Equation (3) and longitudinal data on “domestic independence” available in the IHDS female well-being sample. Panel a. of Table 6 displays the impact on this measure.³⁰ It indicates that the employment guarantee limits domestic independence, decreasing the index by 0.33 (s.e. 0.09)—the variable is standardized to an in-sample mean of 0 and

²⁹This 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 participation and such disagreement generated violence.

³⁰For the results in Table 6, I include individual, age, and spouse age fixed effects as well as controls when specifying ν_i in Equation (3). Appendix Tables A.9 to A.11 are analogous in format to Table 6. They use the other three specifications of ν_i considered throughout the paper and show very similar results.

a standard deviation of 1. For rural married non-disadvantaged women, who are not targeted by the employment guarantee, the impact on the index is smaller. It does not differ statistically from 0 when using standard significance levels.

Though it is noisily estimated, the impact on domestic violence in the other placebo sample, that of urban women, is large in magnitude. There is no reason to expect this large magnitude. Additional evidence indicates that this isolated large impact in the urban sample is not a symptom of concern regarding the identification strategy: the impact on the outcomes below is small in magnitude. It does not differ statistically from 0 when using standard significance levels.³¹

BMI: A Well-Being Measure. Longitudinal data on body-mass index (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 could be inherently subjective and prone to measurement error due to their sensitive content. BMI is an appropriate measure because, in India, it strongly correlates with domestic violence (Ackerson and Subramanian, 2008). It is a measure of mental and physical health (Selvamani and Singh, 2018), which both have been linked to the intra-household distribution of resources (Anderson and Genicot, 2015; Calvi, 2020).

Panel b. of Table 6 indicates that the employment reduces BMI by 0.39 points (s.e. 0.14). Impacts for rural married non-disadvantaged women and urban married women are smaller in magnitude and do not differ statistically from 0 when using standard significance levels. 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. Table 6 verifies that, while it differs from 0 statistically when using standard significance levels, the impact on height for rural married women is small. In the other subsamples, it is small and does not differ statistically from 0. 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).³²

³¹An additional check is in Appendix Table A.8, where I estimate the structural model of resource allocation in the urban sample. The estimates show that the impact on the female share of intra-household consumption is -0.001 (s.e. 0.005).

³²There 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). An average decrease in BMI from the low baseline average of 20.9 in treatment states makes women more vulnerable physically and mentally. A “healthy” decrease in BMI due to a reduction from a high baseline value is not salient in the

Table 6. Female Well-Being and the Employment Guarantee

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
	Panel a. Domestic Independence Index				Panel b. Female Body-Mass Index				Panel c. Female Height in Meters			
	$\bar{y}_{i\text{before}}^{\text{control}}$	$\bar{y}_{i\text{before}}^{\text{treat}}$	ATT	N	$\bar{y}_{i\text{before}}^{\text{control}}$	$\bar{y}_{i\text{before}}^{\text{treat}}$	ATT	N	$\bar{y}_{i\text{before}}^{\text{control}}$	$\bar{y}_{i\text{before}}^{\text{treat}}$	ATT	N
<i>Rural</i>	-0.153	0.083	-0.320 (0.078)	6,388	20.617	20.891	-0.350 (0.122)	8,713	1.507	1.515	0.012 (0.002)	8,713
<i>Rural Non-Disadvantaged</i>	0.106	0.270	-0.076 (0.049)	2,829	22.364	22.615	-0.065 (0.080)	4,018	1.523	1.532	-0.005 (0.002)	4,018
<i>Urban</i>	0.175	0.348	-0.146 (0.079)	3,306	22.716	22.889	-0.011 (0.083)	4,184	1.521	1.522	0.003 (0.002)	4,184

Note: Panel a. displays details from the estimation of the aggregate average treatment on the treated (ATT) based on Equation (3) using the female domestic independence index as the dependent variable. The panel displays ATT estimates for three subsamples (rural, rural non-disadvantaged, and urban). It also displays the corresponding control-state ($\bar{y}_{i\text{before}}^{\text{control}}$) and treatment-state ($\bar{y}_{i\text{before}}^{\text{treat}}$) means of the dependent variable in 2004-2005 (before the employment guarantee) and the number of observations (N). The specification of ν_i includes individual, age, and spouse age fixed effects as well as controls. Panels b. and c. are analogous in format to Panel a. using female body-mass index and height in meters as dependent variables. The standard errors (in parentheses) are bootstrapped clustering at the state \times age-group level. **Sample:** Subsamples of the IHDS female well-being sample indicated in the label.

The evidence in this section has imperfections. It either relies on exclusion and shape restrictions when identifying and estimating structural parameters or on the IHDS female well-being sample, where observations drop due to item non-response (see Section 2). However, its message is cohesive in its diverse sources. While the employment guarantee benefits households as a whole, it hurts the women within them. Precisely, it crowds out their labor force participation, reducing their command of household earnings, intra-household share of consumption, and overall well-being.

7. Comparison to Other Studies

In a literature survey, 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.” Further, he states that identification of mechanisms “demands even more from data and empirical methods.” I aim to fill some of the referenced gaps in the literature by proposing and testing economic frameworks explaining how the employment guarantee shapes work and consumption decisions at the household level as well as within-household resource allocation.

While some studies use the EU-NSS, they do not use all of its available rounds in combination. For instance, while 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 right before the start of the employment guarantee (2004-2005) and one round right 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 recent event-study methods developed for evaluating programs with time-varying roll-out (see de Chaisemartin and D’Haultfoeuille, 2022). Previous studies focus on one data source. I combine the EU-NSS, the HE-NSS, and the IHDS, corroborating findings based on different data sets and empirical strategies and testing implications of my theoretical framework on a variety of outcomes.

I provide 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 he does not use the state-level variation in treatment intensity. He estimates a basic difference-in-difference model (two-period, two-treatment-status regimes) using rounds 2004-2005 and

sample I analyze (only about ten percent of women are overweight and less than one percent are obese).

2007-2008 of the EU-NSS. I reuse the notation in Equation (3) to write his model:

$$\begin{aligned}
 y_{ig} &= \nu_i + \tau_{2007-2008} \\
 &+ \gamma_{2007-2008} \cdot \mathbf{1}[i \text{ lives in a Phase-1 or Phase-2 district}]_i \cdot \mathbf{1}[g = 2007-2008]_g + \varepsilon_{ig},
 \end{aligned} \tag{9}$$

where $g = 2004-2005$ (before the employment guarantee) or $g = 2007-2008$ (after). In Azam (2011), ν_i represents district fixed effects. In his empirical strategy, individuals who reside in Phase-1 or Phase-2 districts are the treatment group; individuals who reside in Phase-3 districts are the control group. He argues that this strategy is plausible because the employment guarantee was not in place in 2004-2005. In 2007-2008, he argues, it was in place only in Phase-1 and Phase-2 districts.³³ He uses a sample of rural women of any marital status who were between 18 and 60 years old at the time of the survey.

I first replicate the estimate of $\gamma_{2007-2008}$ in Azam (2011) using the EU-NSS labor-market sample described in Section 2. I delimit the sample to the rounds of the EU-NSS and age profile that he uses. The details from the estimation are in Column (1) of Table 7. I obtain a point estimate identical to his. The estimation does not include national representativity weights. Column (2) is identical to Column (1) except that it uses national representativity weights. The point estimate halves to 1.2. Its standard error grows and I cannot reject the null hypothesis that it is 0 when using standard significance levels. Column (3) shows that focusing on married women barely changes the point estimate of $\gamma_{2007-2008}$.

The empirical strategy in Azam (2011) can only identify the short-term impact of the employment guarantee; its control group ends up being treated after 2007-2008. An impact estimate of 0 should be obtained when changing the after-treatment period in Equation (9) from 2007-2008 to 2011-2012. The null hypothesis $\gamma_{2011-2012} = 0$ should hold when estimating

$$\begin{aligned}
 y_{ig} &= \nu_i + \tau_{2011-2012} \\
 &+ \gamma_{2011-2012} \cdot \mathbf{1}[i \text{ lives in a Phase-1 or Phase-2 district}]_i \cdot \mathbf{1}[g = 2011-2012]_g \\
 &+ \tilde{\gamma}_{2011-2012} \cdot \mathbf{1}[i \text{ lives in treatment state}]_i \cdot \mathbf{1}[g = 2011-2012]_g + \varepsilon_{ig},
 \end{aligned} \tag{10}$$

while imposing the null hypothesis $\tilde{\gamma}_{2011-2012} = 0$. I present the corresponding estimates of $\gamma_{2011-2012}$ in Columns (4) and (5). These columns only differ in that the former uses the controls in Azam (2011) and the latter uses the controls in this paper. I reject the null hypothesis $\gamma_{2011-2012} = 0$ and thus bring in the IHDS labor-market sample for additional

³³This argument has a caveat. Some of the control-group individuals surveyed in 2007-2008 were already potentially affected by the employment guarantee. In the EU-NSS labor-market sample, 14.8% of the households in Phase-3 districts were surveyed in May of 2008 or later in 2008 in the round 2007-2008.

Table 7. Empirical Comparison to a Common Strategy in Previous Studies

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
	Female Labor Force Participation										log Daily Wage (Males)
$\gamma_{2007-2008}$	0.024 (0.014)	0.012 (0.015)	0.013 (0.017)								
$\gamma_{2011-2012}$				-0.048 (0.019)	-0.033 (0.009)	-0.005 (0.014)	-0.006 (0.013)	-0.006 (0.013)	-0.002 (0.015)		
$\tilde{\gamma}_{2011-2012}$								-0.022 (0.014)	-0.021 (0.014)	-0.042 (0.020)	-0.014 (0.043)
Observations	202,009	202,009	164,630	125,021	131,547	15,758	15,758	15,758	13,777	13,140	14,639
Weights	×	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Individuals FEs	×	×	×	×	×	×	✓	✓	✓	✓	✓
District FEs	✓	✓	✓	✓	✓	✓	×	×	×	×	×
Age FEs	×	×	×	×	✓	✓	✓	✓	✓	✓	✓
Controls, Literature	✓	✓	✓	✓	×	×	×	×	×	×	×
Controls, This Paper	×	×	×	×	✓	✓	✓	✓	✓	✓	✓
Sample States	All All	All All	Married All	Married All	Married All	Married All	Married All	Married All	Married Subset*	Married All	Married All
Age Range	18 to 60	18 to 60	18 to 60	18 to 60	18 to 60	18 to 60	18 to 60	18 to 60	18 to 60	25 to 64	25 to 64
Data Set	EU-NSS	EU-NSS	EU-NSS	EU-NSS	EU-NSS	IHDS	IHDS	IHDS	IHDS	IHDS	IHDS
Years in Sample	2004-05 & 2007-08	2004-05 & 2007-08	2004-05 & 2007-08	2004-05 & 2011-12	2004-05 & 2011-12	2004-05 & 2011-12	2004-05 & 2011-12	2004-05 & 2011-12	2004-05 & 2011-12	2004-05 & 2011-12	2004-05 & 2011-12

Note: Column (1) displays details from the estimation of the aggregate average treatment on the treated for female labor force participation based on Equation (9). The number of observations, use of weights, specification of ν_i , sample, age range of individuals in the sample, data set, and calendar-year coverage are indicated in the rows. Columns (2) and (3) are analogous in format to Column (1). They differ in the details as indicated in the table. Columns (4) to (10) are analogous in format to Column (1). They are based on Equation (10). Coefficient estimates left blank are set to 0. In Column (9), observations from the state of Maharashtra and small territories are not considered. Column (11) is analogous in format to Column (10) for the log daily wage of rural married men. Controls, Literature: controls in Azam (2011) (literacy, caste, age, age squared). Controls, This Paper: controls used throughout this paper. The standard errors (in parentheses) are clustered at the district level in Columns (1) to (9) to make the comparison consistent with the literature. The results and standard errors are directly reproduced from Table 2 and footnote 34 for Columns (10) and (11). **Sample:** Subsample of rural women of the EU-NSS and IHDS labor-market sample (female labor force participation) and subsample of rural women and men of the IHDS labor-market sample.

exploration. Columns (6) and (7) present estimates of this same specification based on this sample. I fail to reject the null hypothesis $\gamma_{2011-2012} = 0$. I conclude that there is no overall consistent support for rejecting this hypothesis.

I then consider estimating Equation (10) without restricting $\gamma_{2011-2012}$ or $\tilde{\gamma}_{2011-2012}$. Such a specification nests my specification of the treatment and control groups and the specification of the treatment and controls groups in Azam (2011). Only the coefficient associated with my specification of the treatment and control groups should differ from 0 when considering a longer time span. Otherwise, it could be that the estimates in Section 4 spuriously pick up a relationship between female labor force participation and the employment guarantee that is wiped out when accounting for the treatment-control specification in Azam (2011). To be clear, the model in Equation (3) is equivalent to the model in Equation (10) when imposing $\gamma_{2011-2012} = 0$. Therefore, $\tilde{\gamma}_{2011-2012}$ is one of the two estimators of the employment-guarantee impact that I use throughout the paper. Column (8) presents estimates of this specification. The estimate of $\gamma_{2011-2012}$ is essentially 0. The estimate of $\tilde{\gamma}_{2011-2012}$ is qualitatively consistent with the evidence in Section 4. The same holds true in Column (9), where I drop observations from the state of Maharashtra and relatively small territories. Azam (2011) suggests dropping these observations given the pre-existence of employment-guarantee programs in the former state and a small number of observations for the latter territories. Column (10) shows that, once delimiting the sample to the age range that I use throughout the paper and imposing $\gamma_{2011-2012} = 0$, the estimate of $\tilde{\gamma}_{2011-2012}$ grows in magnitude and precision.

I rule out that the difference between the results in Azam (2011) and this paper is driven by sample composition (e.g., a specific age profile, marital status, or state of residence) or specification of controls. The difference is due to 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). Indeed, their strategies use the majority of rural Indian districts, those in Phase 3, as the control group. By construction, the estimates of these authors do not contain the impact on Phase-3 districts or the longer-term dynamics driving treatment effects.

I now turn to analyzing the impact of the employment guarantee 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, determining the size and magnitude of such an impact is relevant for my economic interpretation of the employment guarantee in Section 5. Analyzing this impact is not straightforward; it requires addressing two sources of selection. First, selection into the labor force, which is major in the case of women, given their low level of participation. Second, selection into a specific type of work, which is required for observing the corresponding wages. Imbert and Papp (2015) do not consider either source of selection. Their main result is based on the identification strategy in Azam (2011). Their initial sample consists of 356,636 women and men who report either being employed in public or private works (including casual work), unemployed, or out of the labor force. They observe casual-work daily wages (daily earnings from casual work) for 64,167 individuals. This subset of individuals composes the subsample for their analysis of casual-work wages.

I circumvent the first selection issue by only considering rural married men, the majority of men in India. Most of them work. Selectively observing their wages is a secondary concern. Indeed, Imbert and Papp (2015) document that the impact that they find on wages is driven by male wages. I circumvent the second selection issue by analyzing the wage across all working activities, which is essentially observed for all of them. If the employment guarantee has an economically and statistically significant impact on casual-work wages, this should translate into an impact on overall rural wages. I should be able to detect it. Column (11), which uses the same strategy as Column (10), summarizes the results from my analysis. I find a relatively small negative impact of -1.4% (s.e. 4.3%) from a baseline of 5.8 (2018 USD, PPP).³⁴ This finding implies that the 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.³⁵ My finding is almost identical to the finding of Zimmermann (2021) who, focusing on men between 18 and 60 years old and using a different identification strategy than mine, finds that the employment guarantee decreases the wage across working activities by -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

³⁴This impact is based on one of the four specifications that I consider for ν_i when estimating Equation (3) throughout the paper (individual and age fixed effects as well as controls). The four specifications yield the following estimates. Individual fixed effects: -0.5% (s.e. 4.2%). Individual and age fixed effects: -0.3% (s.e. 4.1%). Individual and age fixed effects as well as controls: -1.4% (s.e. 4.3%). Individual, age, and spouse age fixed effects as well as controls: -1.5% (s.e. 4.1%).

³⁵Berg et al. (2018) and Klöpper and Oldiges (2022) report similar findings to Imbert and Papp (2015) regarding sector-specific wages. Their strategies identify a short-term impact. They are also subject to the discussed selection caveats.

Table 8. 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
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 ↑
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 (2021)	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:** 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. 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 the 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 (2021): 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.

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 in my analysis. 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 8 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 Klønner and Oldiges (2022), who also find a decrease in agriculture as the main household occupation. Zimmermann (2021) finds a reallocation towards self-employment for men. All these reallocation results are broadly consistent with Figure 5. The positive impacts on household consumption per capita documented by Bose (2017) and Klønner and Oldiges (2022) are also consistent with my findings in Section 5.

8. Final Comments

India's female labor force participation is salient for its low level and recent decrease, which contrasts with increasing trends around the world. Rural married women drive a recent 25-percent countrywide decrease, observed between 2005 and 2012 from a baseline of 40%. I argue that these women supply labor as added workers, insuring household earnings. An improvement in economic conditions increases their time spent in non-market activities. Social norms that establish a family preference for them not to work at all reinforce this increase. I find that this mechanism prevails when the Mahatma Gandhi National Rural Employment Guarantee Act insures household earnings. A large fraction of women take up employment-guarantee jobs. Yet, the fraction of women who reduce their labor-market activities is even larger. The insurance provided by the employment guarantee increases household consumption, and, therefore, reduces absolute poverty at the household level; it generates richer households. However, it also crowds out female labor force participation, reducing women's command of household earnings. This reduction increases women's intra-household poverty, and thereby, detracts their 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. 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*.
- 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.

- 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.
- 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>.
- de Chaisemartin, C. and X. d’Haultfoeuille (2020). Two-way Fixed Effects Estimators with Heterogeneous Treatment Effects. *American Economic Review* 110(9), 2964–96.
- de Chaisemartin, C. and X. D’Haultfoeuille (2022). Two-Way Fixed Effects and Differences-in-Differences with Heterogeneous Treatment Effects: A Survey. Technical Report w29691, National Bureau of Economic Research, Cambridge, MA.
- 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.
- 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*.
- Goldin, C. and L. F. Katz (2002). The Power of the Pill: Oral Contraceptives and Women's Career and Marriage Decisions. *Journal of Political Economy* 110(4), 730–770.
- 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.
- Kabeer, N. (2008). Paid Work, Women's Empowerment and Gender Justice: Critical Pathways of Social Change.
- 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.
- Ministry of Rural Development (2005a). The National Rural Employment Guarantee Act. Website, https://nrega.nic.in/amendments_2005_2018.pdf.
- Ministry of Rural Development (2005b). The National Rural Employment Guarantee Act: Operational Guidelines. Website, https://web.archive.org/web/20130414230959/http://nrega.nic.in/Nrega_guidelines.pdf.
- 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.
- 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. (2021). Why Guarantee Employment? Evidence from a Large Indian Public-Works Program. *Economic Development and Cultural Change* (Forthcoming).