

ONLINE DATA & EMPIRICAL APPENDIX
PERVERSE CONSEQUENCES OF WELL-INTENTIONED REGULATION: EVIDENCE FROM
INDIA'S CHILD LABOR BAN

PRASHANT BHARADWAJ[†], LEAH K. LAKDAWALA^{††} & NICHOLAS LI^{†‡}

1. DATA

1.1. **Additional Rounds of the NSS**

The 42nd round of the NSS was collected between July 1986 and June 1987. The 42nd round is unique from the employment rounds in that its focus is on “participation in education” rather than on employment. (There are several other modules in this round but none that focuses on employment.) While there exist some employment data in this round, the nature of the employment questions and sampling frame are somewhat different from the employment round questions and thus the employment variables are not consistent across the two subsets of data. Most notably, in all of the rounds *other than* the 42nd Round, employment information is reported for all children 6 and older. In the 42nd Round employment information is reported only for children who are not currently enrolled in school, regardless of how much time they spend in school. Thus child employment variables in the 42nd round are likely to undercount working children relative to the employment rounds. Even considering these caveats, this round of the NSS is potentially useful for evaluating the short-term impacts of the ban because it encompasses the six months immediately preceding and following the Child Labor Act (enacted in December 1986); thus the changes in child time allocation during this time are unlikely to capture long-term trends due to factors other than the 1986 Act.

For the purpose of robustness checks, we also make use of a second additional dataset comprised of the consumption rounds (Schedule 1) of the NSS. These were conducted only in

[†] DEPARTMENT OF ECONOMICS, UNIVERSITY OF CALIFORNIA, SAN DIEGO ^{††} DEPARTMENT OF ECONOMICS, MICHIGAN STATE UNIVERSITY ^{†‡} DEPARTMENT OF ECONOMICS, UNIVERSITY OF TORONTO
E-mail address: prbharadwaj@ucsd.edu, lkl@msu.edu, nick.li@utoronto.ca.
Date: April 2016.

the years *following* the passage of the 1986 Act. These include rounds 45, 46, 47, 48 and 49 which span the period 1989-1993 (there are no consumption rounds available prior to 1986). The consumption rounds contain somewhat limited information on employment (and also lack some controls, such as religion) but include more detailed information on household expenditure.

1.2. Child-level variables

All measures of child time allocation are based on the child's reported "principal usual activity". Each child is reported as having an exclusive principal activity. The reference period for usual activity is the 365 days prior to the survey. A "current status" is also reported for each child with virtually the same possible classifications as for usual status but with the reference period being the week prior to the survey. Current status is not available in the 42nd round of the NSS. All of the main results for the employment sample reported in the paper are robust to using child time allocation variables based on current status rather than usual status.

Aside from the activities that we study in Tables 2 and 3, the other potential categories for child time in the NSS include "Other" (15.6% in 1983), "Too young to attend school/work/seek work" (80.8%) and a few less prevalent activities (e.g. begging, prostitution, disabled etc. which account for the final 3.6%). Thus the increase in child labor force participation seems to come largely from a decrease in unpaid household activities, other activities and being considered "too young" rather than from schooling. The category "Too young to attend school/work/seek work" is available only for the 38th round (1983) after which it is combined with "Other". Therefore we cannot study the impact of the ban on the classification of children as being "too young."

We classify as children working banned versus non-banned occupations based on the 3-digit NIC codes reported for each employed child (according to the principal usual activity); however for the data on days spent working we have only 1-digit NIC codes to match to activities so the banned versus non-banned occupation classifications are much coarser. These are matched to the list of processes and occupations listed as banned in the 1986 Act as of 1993. As stated in the 1986 Act, all children working in family enterprises are not classified as working in banned industries, regardless of the NIC code. Over time, other processes have been added to the "prohibited list" of regulated industries under the 1986 Act. Relatively few of these changes occur between

1986 and 1994. The majority (and more substantive) of the changes to the “prohibited list” occur after 1994, including the prohibition of child employment in domestic work and dhabas (eateries) which were added in October 2006. Note that the identification of the empirical effects of the ban is based solely on age (or sibling age) and year and not sector, so the results in the paper should not be affected by any changes to the “prohibited list”.

As stated in the paper the NSS modules include information days spent in each activity, though the distributions displayed in Figures OA1 and OA2 make clear that the data contain very little variation over and above the extensive margin. Moreover, the data on days in each activity are likely to be reported with some noise, as virtually all children who report attending school report doing so for 7 days a week, even though India’s standard school week covers only 5 days. For that reason, we focus most of our analysis on the principal usual activity reported for each child, which we believe captures mostly the extensive margin of participation in activities. However, even changes in the principal usual activity could potentially capture in part changes on the intensive margin in the sense that they are based on a child’s *primary* activity. For example, if a child goes from primarily attending school to primarily working, this is coded as going from “Any Economic Activity”=0 to “Any Economic Activity”=1, even if the change in status was due to a shift in days devoted to each activity (for example, from full-time schooling and part-time work to full-time work and part-time schooling). The pre-ban data on days spent engaged in economic activity for those whose principal usual activity is *not* work (Figure OA3) suggest that this is possible, given that there is a non-negligible mass of children working an intermediate number of days. Assuming that the principal usual activity corresponds to the activity that comprises the majority of a child’s time, this means that the usual status of these children could change in response to an increase in work days, though it is worth noting that of all children (aged 10-13) whose principal status is not working in the 1983 round, only 1.2% report spending some time working in the past week. The data also contain very limited information on secondary activities. More specifically, secondary activities are only reported if they are “gainful” and thus are not reported if they include schooling or unpaid activities. In the pre-ban period, 4.4% of all children ages 10-13 who are primarily engaged in non-economic activities such as school report some form of gainful secondary activity.

This again suggests a pool of children whose primary activities could shift from non-work to work due to an increase in work hours resulting in a new primary activity classification. In light of these additional sources of data, it seems reasonable that our results reflect in part an intensive margin effect of the ban.

1.3. Household-level variables

We examine the responses of several household-level measures of welfare in Table 6 of the main paper. We now describe these measures in more detail.

Our first measure, per capita expenditure, is one of the most widely used indicators of household welfare and is used for Indian and global poverty measurement. Note that this is the only household outcome available in the 42nd Round. Our second measure, per capita food expenditure, is highly correlated with per capita expenditure (0.92) but may vary if poor households adjust to shocks by lowering non-food intake more than food intake. Our third measure, caloric intake per capita, is moderately correlated with food expenditure per capita (0.54) but can vary if households adjust to a drop in food expenditure by switching from foods that are more expensive per calorie to foods that are cheaper per calorie. Our fourth measure, the staple share of calories, is proposed by Jensen and Miller (2010) as a measure of household nutritional adequacy in the presence of caloric needs that are unknown or variable across households. Their logic is that if households attach a high disutility to having caloric intake below caloric needs, they will substitute towards the cheapest sources of calories (staples). The staple share is thus an inverse indicator of household welfare, and consistent with this we find a negative correlation (-0.57) with per capita expenditures. Finally, our fifth measure is a household asset index constructed as the principal component of a set of variables that capture the quality and quantity of housing, the type of energy used for cooking and lighting, and the quantity of electricity used (which is likely to be correlated with the number of appliances and durables used by the household). If households adjust to negative income shocks by selling assets, using electric appliances less intensively, or letting the quality of their assets deteriorate then we might expect deterioration in the asset index for households affected by the ban. While the asset index is correlated with per capita expenditures (0.50) and other flow measures of welfare, it uses an entirely different source of variation across households.

Thus while measurement error may be correlated across the other measures, it is unlikely to be correlated across our expenditure/consumption flow measures and our asset measure.

We calculate monthly per capita expenditures using Schedule 1.0 of the National Sample Survey (NSS) of India, which is directly linked to the employment data for all households in our sample (both Schedules 1 and 10 were collected in 1983, 1987-8, and 1993-4). The survey is based on a 30-day recall of household consumption for a detailed list of items. Information is collected on both quantities and expenditures and includes home produced goods (which have expenditures imputed at the farm-gate price). Our per capita expenditure measure excludes rent and taxes but includes all food, alcohol and tobacco, energy, clothing and footwear, service, non-durable and durable expenditure. Real values (expressed in 1982 rupees) are nominal values deflated by the average wholesale price index reported by the Government of India for the respective year. We calculate monthly per capita food expenditures using only the food items from the survey and exclude alcohol. To construct caloric intake at the household level we convert the recorded quantities into calories using the standard caloric conversion factors that have been used for this purpose in the past (Gopalan et al. (1980)). We supplement this with data from other sources in a few cases. For some food items quantities and/or caloric conversion factors are not available so we use the imputation procedure described in detail in the appendix of Li and Eli (2013). We use the “liberal” conversion factor and impute using all food factors rather than the group factors. Our calorie calculations include calories from alcohol. We calculate a “staple share of calories” measure as the ratio of calories from cereals and cereal substitutes to calories from all sources (including alcohol). In our regression analysis we define “1 - staple share of calories” as a positive indicator of welfare for reasons we discuss later. We also construct a household wealth/asset index using data on housing and some proxies for durable ownership. Although the 43rd and 50th NSS consumption modules record ownership of many different household durables, these data are missing for the 38th round so we are forced to rely on proxies such as the source of energy for lighting and cooking and household electricity ownership. To calculate the asset index, we calculate the principal component of the following set of discrete variables – source of cooking energy, source of lighting energy, floor

type, wall/building type, house condition – and the continuous variables covered area and quantity of electricity used.

2. ADDITIONAL ROBUSTNESS CHECKS

2.1. Clustering methods: Wild Cluster Bootstrap with Webb Weights

In Online Appendix Tables OA1 and OA2 we use a wild cluster bootstrapping procedure outlined in Cameron et al. (2008). Specifically, we create a counterfactual distribution of t-statistics under the null hypothesis of no effect of the ban. To do this, we run the regression of our outcome of interest on all of our covariates *except* the interaction “Under14XPost-1986” (our “ban” variable) and then calculate the residuals. For each of 999 iterations, we (i) create a “wild” outcome which is the predicted outcome under the null and the aforementioned residual multiplied by Webb weights to reflect the low number of clusters (Webb (2013)) and then (ii) regress the “wild” outcome on our complete set of covariates including “Under14XPost-1986” and store the resulting t-statistic. This gives us a distribution of the t-statistic for “Under14XPost-1986” under the null hypothesis of no effect. The p-value we calculate draws on the comparison of our observed t-statistic to the null distribution.

2.2. Age ranges and sample selection for sibling results

The sample we use for our sibling-based regressions includes all children who have at least one sibling between the ages of 6 and 17. While this age range for siblings is wider than the (own) age range for the “simple difference” estimates in Table 2 of the main paper, it is necessary for several reasons. First, this age range is relatively symmetric around the “treatment-generating” sibling ages (10-13) which is required to balance the estimating sample between “treatment” and “control” children. If, for example, we asymmetrically restricted the sample to children with siblings ages 10-17, we would be selectively dropping (control) children with younger siblings from our sample. In fact, we find that doing so would yield a sample highly skewed towards “treatment” children (65% versus 35% “control”), whereas the 6-17 sibling range yields a more balanced sample (48.1% “treatment” versus 51.9% “control”).

Second, and more importantly, the symmetric sibling age range is necessary to avoid bias due to sample selection. This is because sibling age is determinant of own work status, even in the absence of the ban (Manacorda (2006)). To see why this causes bias, consider again the example in which we asymmetrically restrict the sample to children with siblings ages 10-17. In this case all “control” children will have older siblings (ages 14-17) and no “control” children will have younger siblings. Since sibling age affects own work status, this means that our sample of “control” children will be selected on having a relatively low likelihood of work (our outcome of interest), as children with older siblings are less likely to work than children with younger siblings. Therefore using an asymmetric sample is effectively selecting our sample in a way that is correlated with our outcome of interest, leading to biased estimates. For both of these reasons we use a symmetric range of sibling ages for our sibling-based estimation. In column 4 of Appendix Table A.1 we show that our results are robust to using an even narrower sample of siblings ages 8-15; however, it should be noted that this narrow sample is less balanced on treatment (59%) and control (41%) children.

2.3. Implied changes in child productivity

We observe a net zero impact of the ban on household expenditure, despite the increase in child work.¹ To back out an implied change in child productivity, we use our additional results that other observable components of household income - adult labor income and assets - also do not change in response to the ban. We do not observe saving, transfers or other non-labor income in the data. Assuming that percentage change in household expenditure and other non-child labor components are zero implies that the increase in child work and decrease in child productivity must exactly offset each other (in percentage terms). To calculate the increase in child labor supply at the household labor we separately estimate the effects on the proportion of working children ages 6-9 and 10-13 at the household level.² The effect for the 6-9 age group is a 25.3% increase over

¹This holds for both the full sample of households and the restricted sample of households with at least one child aged 6-17.

²Labor supply responses at the household level must be estimated separately by age group because the definition of treatment changes as the age group for the outcome changes. Specifically when estimating the effect of the ban on the proportion of children ages 6-9 working, the treatment variable is whether the household has at least 1 child in the age range 10-13. However when the outcome is the proportion of children age 10-13 the treatment is whether the household has at least 2 children ages 10-13 because all children in the sample have at least 1 child age 10-13

the pre-ban mean and for the 0-13 age group it is 6.54% (results available upon request). We then calculate the proportion of child labor derived from each age group to use as weights (0.115 and 0.885) for ages 6-9 and 10-13, respectively. The overall effect on child labor supply at the household level is then the weighted average of the effects for the 6-9 and the 10-13 age groups:

$$0.8846 * 0.0654 + (1 - 0.8846) * 0.2532 = 0.0871.$$

REFERENCES

- Cameron, A. C., J. B. Gelbach, and D. L. Miller (2008). Bootstrap-based improvements for inference with clustered errors. *The Review of Economics and Statistics* 90(3), 414–427.
- Gopalan, C., B. Rama Sastri, and S. Balasubramanian (1980). *Nutritive Values of Indian Foods (3rd edition)*. Hyderabad: National Institute of Nutrition, Indian Council of Medical Research.
- Jensen, R. T. and N. H. Miller (2010). A revealed preference approach to measuring hunger and undernutrition. *NBER Working Paper* (16555).
- Li, N. and S. Eli (2013). Can caloric needs explain three food consumption puzzles? evidence from india. *Working Paper*.
- Manacorda, M. (2006). Child labor and the labor supply of other household members: Evidence from 1920 america. *The American Economic Review* 96(5), 1788–1801.
- Webb, M. D. (2013). Reworking wild bootstrap based inference for clustered errors. Technical report, Queen’s Economics Department Working Paper.

(otherwise the outcome variable is undefined). These definitions are consistent with the sibling-based definition of treatment.

3. ONLINE APPENDIX FIGURES AND TABLES

FIGURE OA1. Distribution of days spent in any economic activity in previous week (Ages 10-17, 1983).

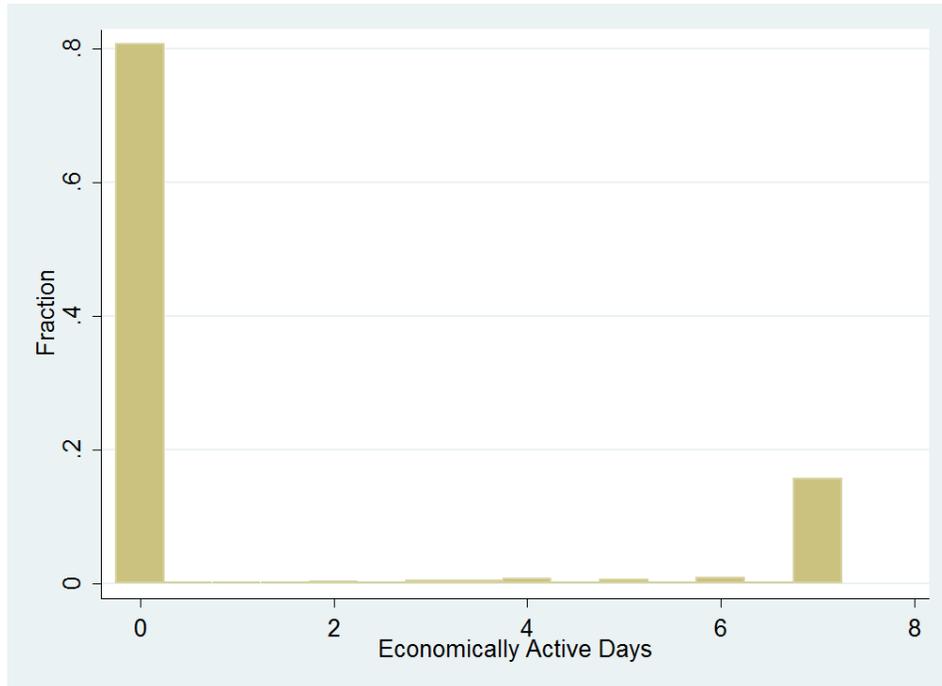


FIGURE OA2. Distribution of days spent attending school in previous week (Ages 10-17, 1983).

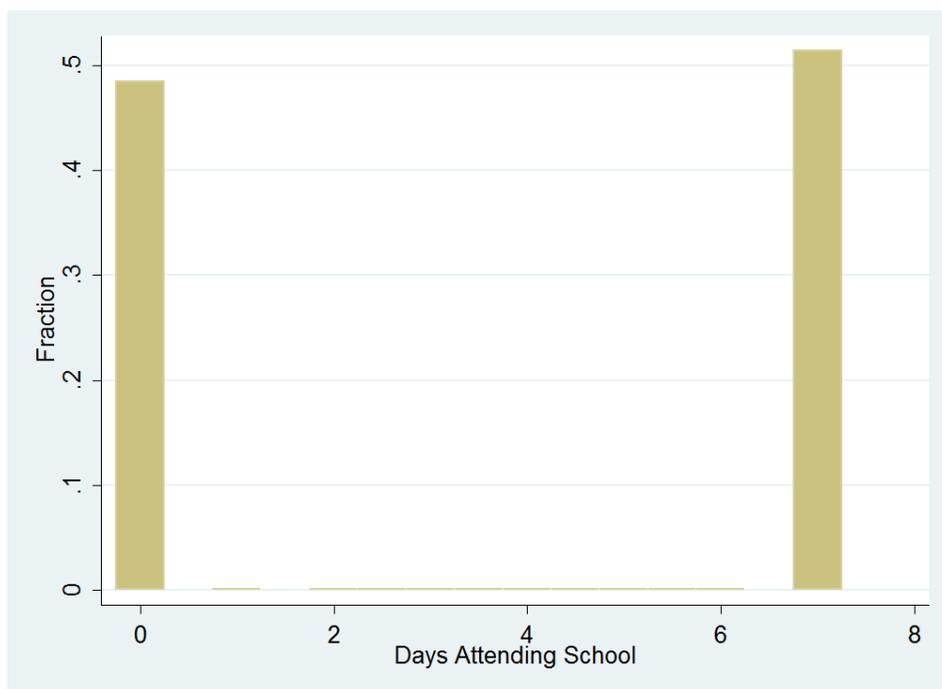
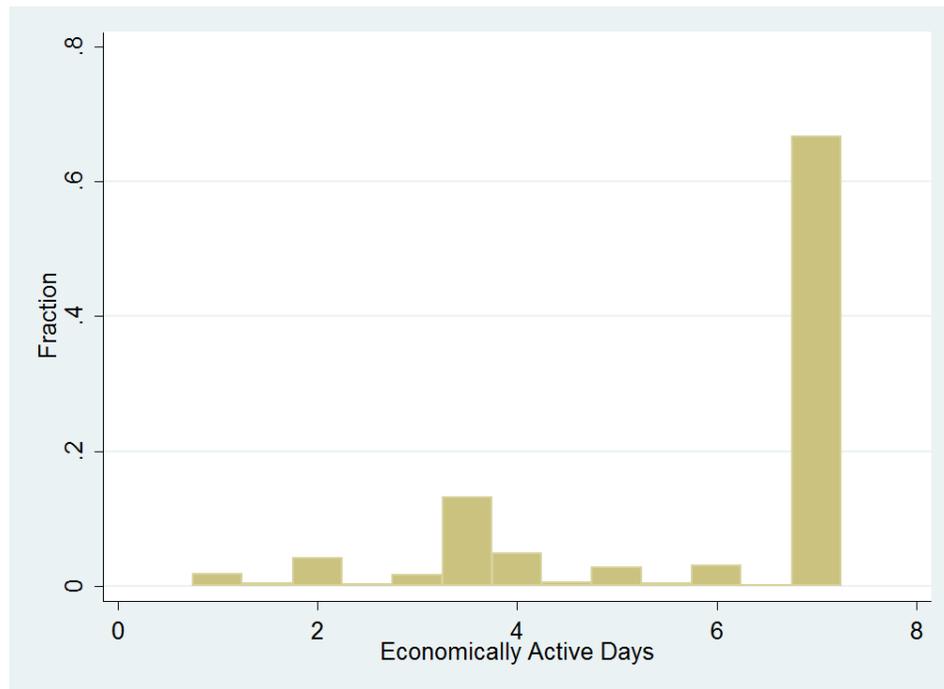


FIGURE OA3. Distribution of days spent in any economic activity in previous week, conditional on principal usual status *not* being economic activity (Ages 10-13, 1983).



This distribution of days worked in the past *week* is for the sample of children ages 10-13 in 1983 whose principal usual status (based on a recall period of the past *year*) is *not* engaging in economic activity but who have spent at least some time in economic activity in the past week. Note that only 1.2% of children ages 10-13 in 1983 whose principal usual status is *not* working report spending some time working in the past week.

FIGURE OA4. Raw probability of economic activity by age, pre- and post-1986.

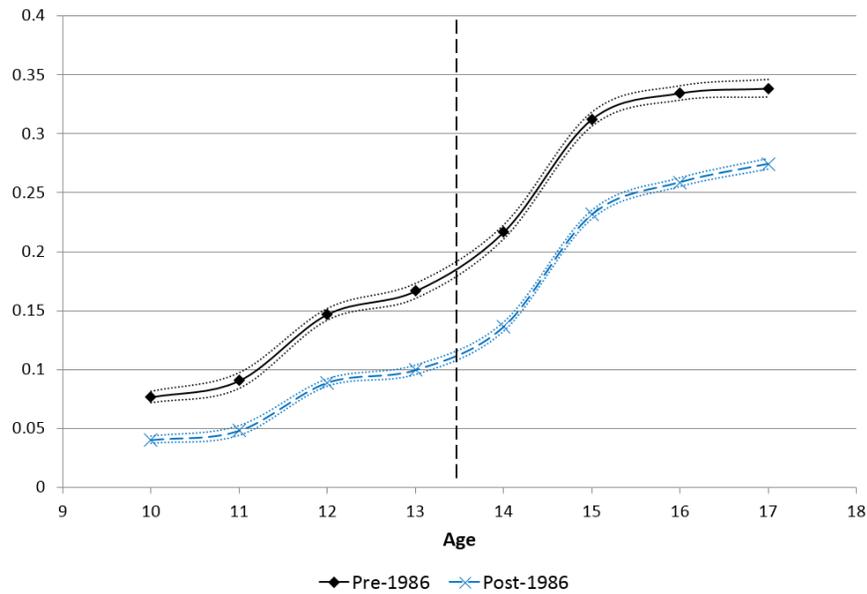
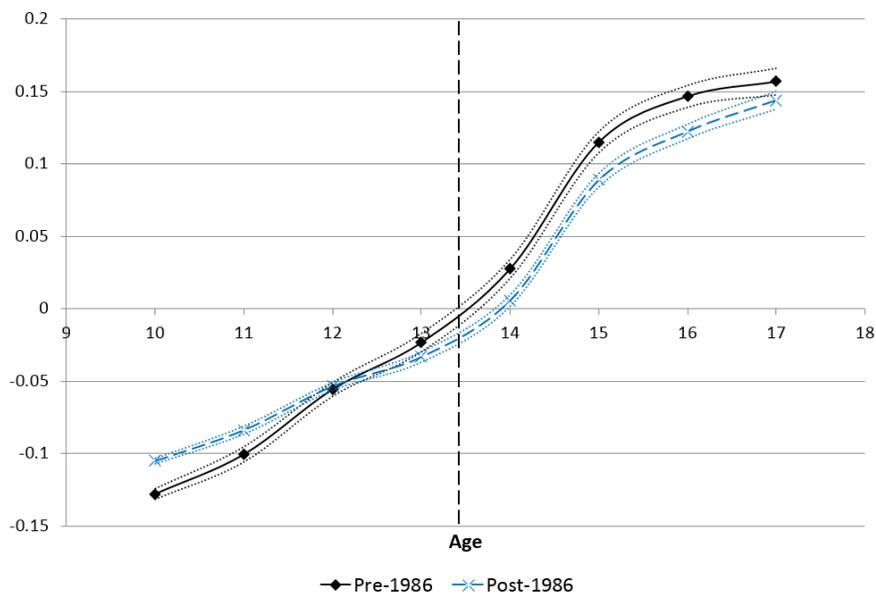


FIGURE OA5. Residual probability of economic activity by age, pre- and post-1986.



Residual probabilities are probabilities of economic activity after partialling out the effect of the following variables: gender, gender of household head, age of household head, urban status, number of adult females, number of male children, number of female children, number of children under 5, number of children ages 6-9 as well as the following fixed effects: family size, household head's education level, religion, survey quarter, state.

FIGURE OA6. Raw changes in the probability of economic activity (post ban - pre ban) by age and sibling-based treatment status.

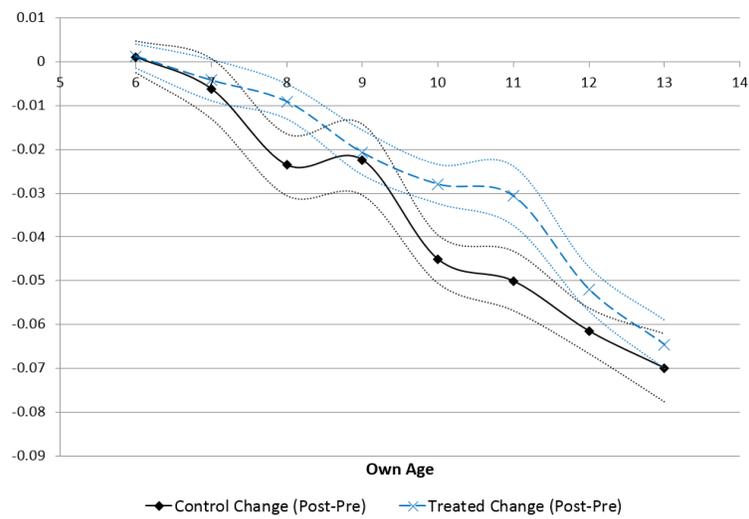
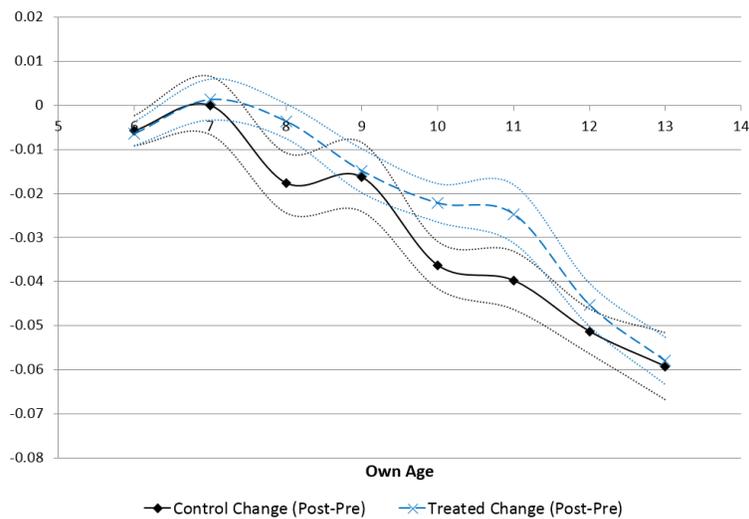
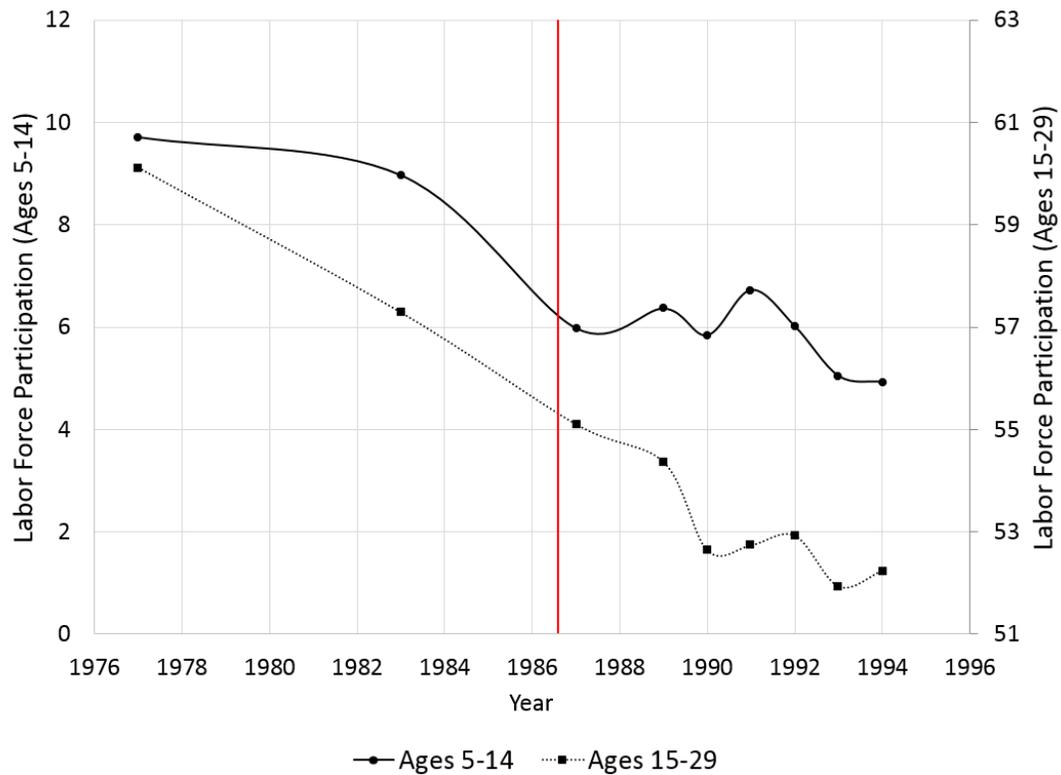


FIGURE OA7. Residual changes in the probability of economic activity (post ban - pre ban) by age and sibling-based treatment status.



Residual probabilities are probabilities of economic activity after partialling out the effect of the following variables: gender, gender of household head, age of household head, urban status, number of adult females, number of male children, number of female children, number of children under 5, number of children ages 6-9; a separate variable for the number of household members of each age 0-25; as well as the following fixed effects: family size, household head's education level, religion, survey quarter, state.

FIGURE OA8. National trends in labor by age group (5-14 versus 15-29) from 1977-1994.



Data by age group, sex, and sector (rural, urban) are collected by the Central Statistical Organisation, Ministry of Statistics and Programme Implementation, Govt. of India, and reported by IndiaStat (www.indiastat.com). Age groups are as reported by IndiaStata and are *not* chosen by the authors; note that they do not perfectly align with the age restrictions of the 1986 Act. Additional data from the 1981 and 1991 Censuses are used to aggregate the labor force participation data to the national level; data from the IndiaStat survey is matched to the closest Census in absolute terms. The horizontal axis marks out the first calendar year of each corresponding survey, even if the survey spans two calendar years; for example the first survey was conducted in 1977-78 and is thus marked at 1977 (not 1978). Labor Force Participation of age groups 5-14 and 15-29 are plotted using separate axes with different intercepts but the scale for both axes is the same.

FIGURE OA9. Age distribution pre- and post-ban.

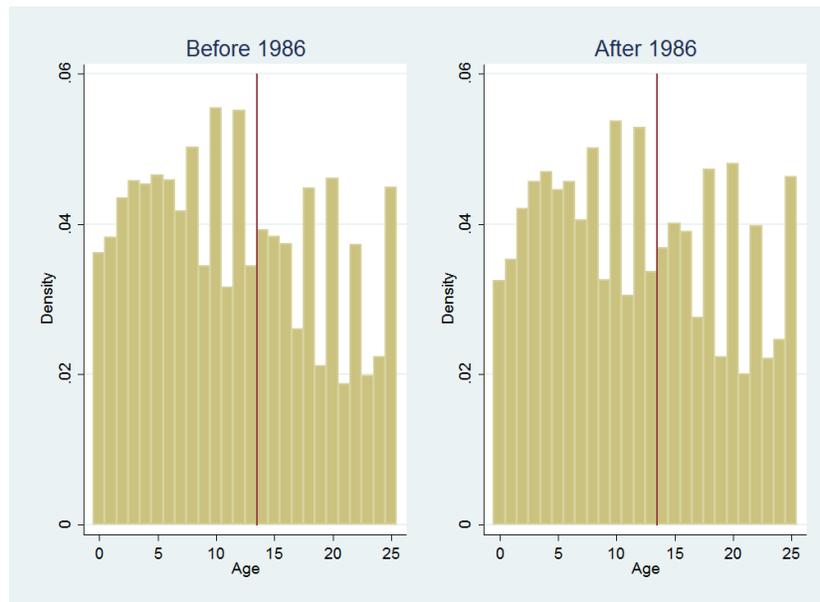


FIGURE OA10. Age distribution pre- and post-ban for children working in banned occupations.

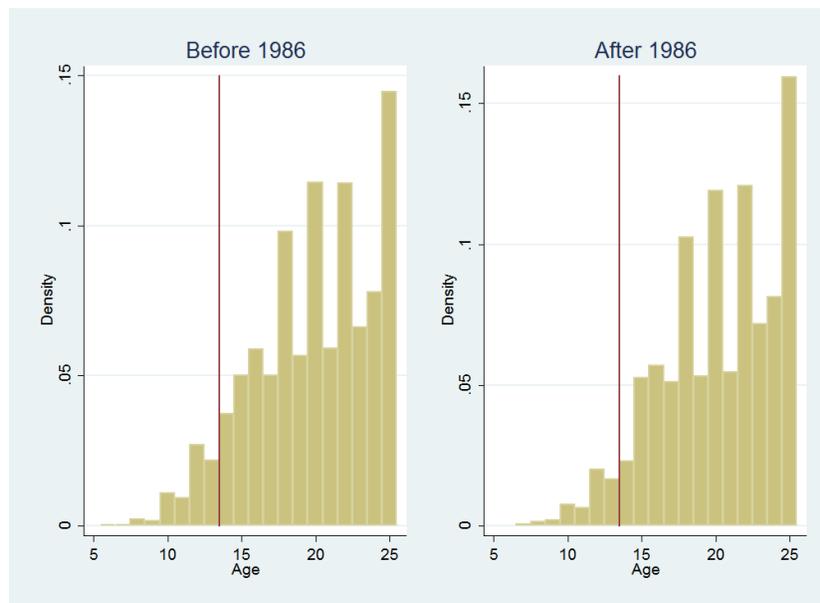


TABLE OA1. Alternate Clustering Methods

Dependent Variable: Any Economic Activity						
	Standard Cluster by Age-Round (1)	Standard Cluster by State (2)	Standard Cluster by Age (3)	Wild Cluster Bootstrap by Age (4)	Standard Cluster by Under 14-Post (5)	Wild Cluster Bootstrap by Under 14-Post (6)
Under14XPost	0.026*** (0.005)	0.026*** (0.005)	0.026*** (0.007)	0.026** N/A	0.026*** (0.000)	0.026 N/A
No. of clusters	24	31	8	8	4	4
p-value	0.000	0.000	0.005	0.010	0.000	0.176
Observations	327,233	327,233	327,233	327,233	327,233	327,233

Columns (4) and (6): Wild cluster bootstrap is implemented as in Cameron, Gelbach and Miller (2008) but using the 6-point distribution weights presented in Webb (2012) due to the low number of clusters. See the Online Appendix for full details on bootstrapping methods. “Under14” is a dummy variable that takes the value of 1 if the child is under 14 years of age. Sample consists of all individuals related to the household head aged 10-17. Controls: gender, gender of household head, age of household head, urban status, number of adult females, number of male children, number of female children, number of children under 5, number of children ages 6-9 as well as the following fixed effects: age, family size, household head’s education level, religion, survey round, survey quarter, state.

TABLE OA2. Simple Estimates of the Effects of the Ban - Narrower Age Ranges

	Dependent Variable: Any Economic Activity			
	Ages 10-17 (1)	Ages 11-16 (2)	Ages 12-15 (3)	Ages 13-14 (4)
Under14XPost	0.026***	0.024***	0.019**	0.011
CRVE (age-round)	(0.005)	(0.004)	(0.004)	(0.003)
Bootstrap p-value	0.000	0.006	0.030	0.340
Number of Clusters	24	18	12	6
Pre-Ban Mean	0.118	0.138	0.154	0.167
Observations	327,233	241,301	169,995	72,964
R-squared	0.182	0.177	0.160	0.136

*** p<0.01, ** p<0.05, * p<0.1 based on the p-value using the Wild Cluster Bootstrap method. "Under 14" is a dummy variable that takes the value of 1 if the child is under 14 years of age. Controls: gender, gender of household head, age of household head, urban status, number of adult females, number of male children, number of female children, number of children under 5, number of children ages 6-9 as well as the following fixed effects: age, family size, household head's education level, religion, survey round, survey quarter, state. Sample consists of all individuals related to the household head aged 10-17. CRVE: Standard errors are given by the conventional cluster-robust estimate of the variance matrix, where the cluster level is age-survey round. Bootstrap p-values are calculated using the Wild cluster bootstrap method (Cameron, Gelbach and Miller (2008)) using the 6-point distribution weights presented in Webb (2012). Pre-Ban mean is for children under the age of 14 only.

TABLE OA3. Simple Estimates of the Effects of the Ban
Round 42: July 1986 - June 1987

	Any Economic Activity (1)	Any Economic Activity (2)	Unpaid Economic Activity (3)	Paid Employment (4)	Enrolled in School (5)	Unpaid Household Services (6)
Under14XPost	0.004 (0.029)	0.003 (0.004)	-0.002 (0.004)	0.006** (0.002)	-0.017*** (0.004)	-0.001 (0.004)
Pre-Ban Mean of Dep. Var.	0.059	0.059	0.034	0.024	0.743	0.093
Observations	90,248	90,248	90,248	90,248	90,248	90,248
R-squared	0.045	0.141	0.086	0.082	0.248	0.212
Controls?	No	Yes	Yes	Yes	Yes	Yes

*** p<0.01, ** p<0.05, * p<0.1 "Under 14" is a dummy variable that takes the value of 1 if the child is under 14 years of age. Controls: gender, gender and age of HH; urban status; counts of adult females, male children, female children, children under 5, children ages 6-9; and the following fixed effects: age, family size, HH head's educ., survey quarter, district. Pre-Ban mean is for children under 14 only. Standard errors are clustered by age-quarter.

TABLE OA4. Sibling-based Estimates of the Effects of the Ban on Child Activities in Previous Week (Days)

Panel A: Children Ages 6-9							
	Any Economic Activity (1)	Employment in Banned Occ. (2)	Employment in Non-Banned Occ. (3)	Unpaid Econ. Activity (4)	Paid Employment (5)	Attending School (6)	Unpaid Household Services (7)
SibUnder14XPost	0.041*** (0.011)	0.001 (0.003)	0.040*** (0.011)	0.040*** (0.010)	-0.001 (0.005)	-0.060 (0.041)	0.016 (0.012)
Pre-Ban Mean of Dep. Var.	0.112	0.005	0.112	0.091	0.019	4.026	0.161
Observations	152,882	152,882	152,882	152,882	152,882	152,882	152,882
R-squared	0.025	0.003	0.025	0.021	0.007	0.260	0.024
Panel B: Children Ages 10-13							
	Any Economic Activity (1)	Employment in Banned Occ. (2)	Employment in Non-Banned Occ. (3)	Unpaid Econ. Activity (4)	Paid Employment (5)	Attending School (6)	Unpaid Household Services (7)
SibUnder14XPost	0.054*** (0.021)	0.012* (0.007)	0.054*** (0.021)	0.041** (0.018)	0.013 (0.012)	-0.064* (0.033)	0.000 (0.022)
Pre-Ban Mean of Dep. Var.	0.770	0.105	0.671	0.524	0.232	4.155	0.900
Observations	164,200	164,200	164,200	164,200	164,200	164,200	164,200
R-squared	0.096	0.015	0.097	0.064	0.045	0.257	0.132

*** p<0.01, ** p<0.05, * p<0.1 “SibUnder14” is a dummy variable for whether the child has at least 1 sibling age 10-13.

Controls: gender, gender of HH head, age of HH head, urban status, no. of adult females, no. of male children, no. of female children, no. of children under 5, no. of children ages 6-9 and the following fixed effects: (own) age, family size, HH head’s education level, religion, survey round, survey quarter, state. Additionally for each age 0-25, we include a separate variable which counts the number HH members of that specific age. Sample consists of all individuals related to the HH head with at least one other (related) HH member age 6-17. Standard errors are clustered by HH. Note that classification of occupations as banned or non-banned uses 1-digit NIC codes (different than for principal usual activity).

TABLE OA5. Household Heterogeneity

Dependent Variable: Any Economic Activity						
	HH Head Has Less than Sec. Educ		Non-Staple Share of Calories		Scheduled Caste Status	
	Overall Effect (1)	Sibling Effect (2)	Overall Effect (3)	Sibling Effect (4)	Overall Effect (5)	Sibling Effect (6)
Heterog.XUnder14XPost	0.016*** (0.005)	0.011** (0.005)	0.007* (0.004)	0.007 (0.006)	0.024** (0.010)	0.003 (0.012)
Under14XPost	0.003 (0.002)	-0.001 (0.003)	0.015*** (0.004)	0.004 (0.004)	0.025*** (0.004)	0.009*** (0.003)
Age Group	10-17	10-13	10-17	10-13	10-17	10-13
Observations	326,754	175,374	318,570	170,920	327,223	175,618
R-squared	0.197	0.106	0.198	0.110	0.192	0.109

*** p<0.01, ** p<0.05, * p<0.1. Heterogeneity Measures: Cols (1)-(2) – Dummy variable for whether the household head has educational achievement less than secondary school; Cols (3)-(4) – “Non-staple share of calories” is a dummy variable for whether the household’s share of daily calories from sources other than cereals and cereal substitutes is *below* the pre-ban (1983) median share; Cols (5)-(6) – “Scheduled Caste” is a dummy variable for whether the household belongs to a scheduled caste. Controls for all columns: gender, gender of household head, age of household head, urban status, number of adult females, number of male children, number of female children, number of children under 5, number of children ages 6-9 as well as the following fixed effects: age, family size, household head’s education level, religion, survey round, survey quarter, state. All regressions also include interactions between all control variables (including “Under14” and “Post”) and the measure of heterogeneity. Columns (1),(3), (5): “Under14” is a dummy variable that takes the value of 1 if the child is under 14 years of age. Sample consists of all individuals related to the household head aged 10-17. Standard errors are clustered by age-survey round. Columns (2),(4),(6): “Under14’ is a dummy variable that takes the value of 1 if the child has at least one sibling age 10-13. In addition to the controls already listed, we include controls for sibling age as follows: for each age 0-25, we create a separate variable which counts the number of household members of that specific age. Sample consists of all individuals related to the household head with at least one other (related) household member age 6-17. Standard errors are clustered by household.

TABLE OA6. Selective Wage Reporting by Industry and Age Group

	Proportion of working individuals reporting wages (%)			
	Banned Industries		Non-Banned Industries	
	Ages 6-13	Ages 14-21	Ages 6-13	Ages 14-21
	(1)	(2)	(3)	(4)
Before 1986	90.79	90.06	24.78	29.90
After 1986	58.03	63.28	14.08	18.42
Difference	-32.76	-26.78	-10.70	-11.48

TABLE OA7. Excluding Post Rounds

	Dependent Variable: Any Economic Activity			
	Excluding Round 43 (1987-8)		Excluding Round 50 (1993-4)	
	Ages 6-9	Ages 10-13	Ages 6-9	Ages 10-13
	(1)	(2)	(3)	(4)
SibUnder14XPost	0.006*** (0.002)	0.007* (0.003)	0.007*** (0.002)	0.009*** (0.003)
Survey Years Included	1983 (pre-ban), 1993-4 (post-ban)		1983 (pre-ban), 1987-8 (post-ban)	
Observations	97,420	105,295	109,753	116,168
R-squared	0.027	0.109	0.027	0.110

*** p<0.01, ** p<0.05, * p<0.1 Columns (1) and (3): “Under14”=dummy for whether child is under 14. Sample: all related to the HH head aged 10-17. SEs clustered by age-survey round. Columns (2) and (4): “Under14”=dummy for whether the child has at least one sibling age 10-13. Sample: all related to the HH head with at least one other (related) HH member age 6-17. SEs clustered by HH. Controls: gender, gender and age of HH; urban status; counts of adult females, male children, female children, children under 5, children ages 6-9; and the following fixed effects: age, family size, HH head’s educ., relig., survey round, survey quarter, state. Columns (2) and (4): for each age 0-25, we include a separate variable which counts the no. of household members of that specific age.

TABLE OA8. Effects on Other Ages

	Dependent Variable: Any Economic Activity					
	Ages 14-17 (1)	Ages 18-25 (2)	Ages 26-35 (3)	Ages 36-45 (4)	Ages 46-55 (5)	Ages 56+ (6)
ChildUnder14	-0.004	-0.002	0.000	-0.010***	-0.001	0.003
XPost	(0.012)	(0.004)	(0.003)	(0.004)	(0.005)	(0.006)
Mean of Dep. Var.	0.296	0.520	0.608	0.677	0.664	0.385
Observations	145,389	173,481	189,540	166,627	92,837	84,433
R-squared	0.196	0.337	0.488	0.516	0.508	0.423

*** p<0.01, ** p<0.05, * p<0.1 “ChildUnder14” is a dummy variable that takes the value of 1 if there is at least one child age 10-13 in the household. Controls: gender, gender of household head, age of household head, urban status, number of adult females, number of male children, number of female children, number of children under 5, number of children ages 6-9 as well as the following fixed effects: (own) age, family size, household head’s education level, religion, survey round, survey quarter, state. Additionally we include controls for sibling age as follows: for *each* age 0-25, we create a separate variable which counts the number household members of that specific age. Sample consists of all individuals related to the household head living with at least one (related) household member age 6-17. Standard errors are clustered by household.