

NBER WORKING PAPER SERIES

EXPLAINING THE INCOME AND CONSUMPTION EFFECTS OF COVID IN INDIA

Arpit Gupta
Anup Malani
Bartek Woda

Working Paper 28935
<http://www.nber.org/papers/w28935>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
June 2021

We thank Mushfiq Mobarak and Cynthia Kinnan (discussants), as well as seminar participants at the IIM Calcutta-NYU Stern India Research Conference, CAFRAL, University of Chicago, and the CPHS Research Seminar for helpful comments. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2021 by Arpit Gupta, Anup Malani, and Bartek Woda. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Explaining the Income and Consumption Effects of COVID in India
Arpit Gupta, Anup Malani, and Bartek Woda
NBER Working Paper No. 28935
June 2021
JEL No. E2,O1,O53

ABSTRACT

The COVID-19 pandemic led to stark reductions in economic activity in India. We employ CMIE's Consumer Pyramids Household Survey to examine the timing, distribution, and mechanism of the impacts from this shock on income and consumption through December 2020. First, we estimate large and heterogeneous drops in income, with ambiguous effects on inequality. While incomes of salaried workers fell 35%; incomes of daily laborers fell 75%. At the same time, we observe that income fell more for individuals from households in the highest income quartile. Second, we document an increase in effort to buffer income shocks by switching occupations. We employ a Roy Model to estimate the gains from occupation churn and find, surprisingly, that reservation wages fell, implying that the risk of COVID did not reduce the value of employment. Third, we find that consumption fell less than income, suggesting households were able to smooth the idiosyncratic components of the COVID shock as well as they did before COVID. Finally, consumption of food and fuel fell less than consumption of durables such as clothing and appliances. Following Costa (2001) and Hamilton (2001), we estimate Engel curves and find that changes in consumption reflect large price shocks (rather than a retreat to subsistence) in sectors other than food and fuel/power. In the food sector, it appears that lockdown successfully distinguished essential and non-essential services, at least to the extent that it did not increase the relative price of food. There is some suggestive evidence that the price shocks outside the food sector were larger in places with greater COVID-19 cases, even during the lockdown.

Arpit Gupta
Stern School of Business
44 West 4th Street
New York, NY 10012
agupta3@stern.nyu.edu

Bartek Woda
University of Chicago
Law School
1111 E. 60th St.
Chicago, IL 60637
woda@uchicago.edu

Anup Malani
University of Chicago Law School
1111 E. 60th Street
Chicago, IL 60637
and NBER
amalani@uchicago.edu

1 Introduction

India was hit quickly and hard by SARS-CoV-2 (Figure 1A). Although the first case was reported only on January 30, 2020, urban slums and some states saw a 50% estimated seroprevalence by July–August 2020. (Mohanani et al., 2020; Malani et al., 2021), leading to the world’s second-highest number of reported cases.¹ Even reported cases are likely an underestimate: seroprevalence studies suggest that official reports may underestimate cases by 30–100 times (Mohanani et al., 2020; Malani et al., 2021).

The economic shocks resulting from the pandemic itself were compounded by a severe series of suppression policies (Figure 1B). India gradually increased international and local travel bans leading to a nation-wide lockdown on 24 March, 2020. This lockdown was among the harshest in the world as measured either by ratings in the Oxford Tracker (Hale et al., 2020) or by Google mobility trends (Google LLC, 2021).² These joint shocks of pandemic severity and government responses make it particularly important to study the economic responses in the context of a developing economy.

We examine the combined effects of the pandemic and resulting restrictions on economic outcomes as measured by the Centre for Monitoring Indian Economy’s (CMIE) Consumer Pyramids Household Survey. Our paper has both a descriptive component and an analytic one.

Our descriptive analysis first documents the decline in income and consumption associated with COVID-19. We document the timing, extent and distribution of the decline (relative to 2019 averages) in activity. This analysis shows that employment began to decline before lockdown, suggesting that voluntary social distancing played a role in the economic effects of COVID-19. Although the decline was most dramatic during lockdown, there was a quick V-shaped recovery such that activity was down only 15% on average by December 2020.

Thereafter, our descriptive analysis demonstrates that the economic shocks had a more heterogeneous effect on income than consumption.³ While salaried workers saw mean income decline by 35% at the worst point in April 2020, daily workers saw their mean income decline a staggering 75%. However, the consumption response was not very unequal: all occupational groups saw mean consumption fall roughly 40%. (We find similar patterns when we differentiate populations based on pre-COVID income rather than occupation.) These findings are consistent with credit markets, or some other insurance mechanism, helping to buffer income shocks during the crisis.

It should be noted, however, that COVID may not have worsened income inequality. At the same time that lower wage occupations seem to suffer a greater decline in income, individuals from higher income households in rural areas experience a bigger decline in


¹As of March 2021 (Worldometer, 2021).

²We do not contend that the decline in mobility is entirely caused by government orders. In fact, mobility fell even before the national lockdown. However, to the extent that the lockdown did reduce mobility, mobility is measure of the severity of suppression.


³This finding is consistent with the idea that inequality in income does not always translate into inequality in consumption Aguiar and Bils (2015); Krueger and Perri (2006) This analysis is similar to Chaudhuri and Paxson (2002), which finds that seasonal variation in consumption is less severe than such variation in income in rural India.

Figure 1: COVID Trajectory, Severity of Lockdown, and Mobility Changes

Panel A: Cases and deaths



Panel B: Lockdown severity and mobility



Note. Case and death data are from www.covid19India.org. We show aggregated daily reported cases and deaths from the government. Shaded period marks the national lockdown. Lockdown severity data are from Oxford Hale et al. (2020). Mobility data are from Google mobility reports Google LLC (2021). Shaded period marks the national lockdown. Time periods cover February 2020–January 2021.

income than those from higher income households. (In urban areas, individual incomes decline about the same percentage across the household income distribution.) During the recovery, individuals from higher income households in both rural and urban experience the weakest recovery in incomes. The discrepancy between the occupation and income results for inequality may be potentially explained by household composition. Lower income individuals in higher income households may be the first to lose or sacrifice income and least likely to return to the labor market for fear of contracting COVID.

The core of our paper examines how individuals managed to cope with this large economic shock. With respect to income, we focus on two margins. We first examine whether the decline in economic activity reduced work on the intensive margin of hours versus the extensive margin of employment. We find that COVID reduced male employment by 20 percentage points (from a base of 65%) and reduced male hours worked (conditional on employment) by 1.5 hours per day (from a base of 8 hours).⁴

Second, we document how individuals switched occupations (without migrating) to offset declines in their income. Those who remained employed during the pandemic increased the rate at which they changed occupations by 50% over their 2019 levels of baseline occupation switching. We use a Roy Model to estimate the gain from switching occupation to protect income and find that reservation wages and the gains to switching declined during the pandemic. The decline in reservation wage suggests that fear of COVID is not what increased unemployment: that fear would have raised reservation wages. The decline in reservation wages is also consistent with the idea that aggregate shocks to household consumption reduce consumption in the unemployed state because other household member were also more likely to become unemployed.

With respect to consumption, we do two things. First, we examine the degree of consumption smoothing using Cochrane (1991) or Townsend (1994) style regressions. We find that households were about as successful at smoothing idiosyncratic shocks during COVID as before COVID. While this does not suggest full compensation—as COVID was a large aggregate shock—it does suggest a surprising strength of formal and informal insurance mechanisms.

Second, we examine how household allocated their budget across product categories. We document that individuals reduced their expenditures on some goods (clothes, appliances, education) far more than on other goods (food, fuel and power, and housing and rent). We then estimate Engel curves implied by the Almost Ideal Demand System of Deaton and Muellbauer (1980) to decompose the decline in expenditure into two parts. One is movement along the pre-COVID Engel curve due to declines in income, holding the relative prices of goods constant. The other is declines in expenditure due to increases in the relative price of goods (Hamilton, 2001). This analysis reveals that price shocks for essential goods such as food and fuel were far less severe than price shocks for other, non-essential goods. Food and fuel budget shares increased because obey Engel's law and income declined during COVID. Consumption of categories like clothing suffered mainly because prices, as perceived by consumers, rose. Examining price shocks across product categories and locations we find that India's lockdown seems to have successfully distin-

⁴Female employment fell about 3 percentage points from a base of 9%. Female hours conditional on employment fell less than 1 hour from a base of 7 hours per day.

guished between essential (food) and non-essential goods (clothing). Moreover, we find that the price shock outside food was positively correlated with COVID-19 cases. This may be evidence that the lockdown had a harsher effect on markets where cases were higher or that the shadow price of non-essential goods reflects peoples' fears of catching COVID.

Our work is related to three lines of research. Our analysis of incomes relates to work on income smoothing in developing countries (Morduch, 1995; Dercon, 2002). Prior literature has focused on income smoothing by choice and timing of inputs into farming (Binswanger and Rosenzweig, 1993), occupational diversification within households (Rosenzweig and Stark, 1989), migration (Paulson, 1994), and contracting (Bardhan, 1983). Our work relates more closely to work on smoothing by labor supply on the intensive margin (Moser, 1998) or across sectors (Kochar, 1999; Rose, 2001). Our work also relates closely to the literature on occupational mobility (Moscarini and Thomsson, 2007; Kambourov and Manovskii, 2008, 2009a,b; Groes et al., 2015), how it relates to the business cycle (Carrillo-Tudela and Visschers, 2020), and the collateral consequences of that adaptation (Chodorow-Reich and Wieland, 2020).

Our analysis of consumption relates to a literature that employs Engel curves to identify bias in CPI measurements (Nakamura, 1995; Hamilton, 2001; Costa, 2001; Logan, 2009; Atkin et al., 2020) or changes or inequality in income (Almås, 2012; Young, 2012; Aguiar and Bils, 2015; Nakamura et al., 2016). Our work differs in that our goal is not to identify implicit price changes for the purpose of correcting CPI so much as finding implicit price changes due to the COVID supply shock.

Finally, our analysis contributes to a growing literature that seeks to understand the impact of COVID on economic outcomes overall (Alstadsæter et al., 2020; Chen et al., 2020) and on inequality (Deaton, 2021; Egger et al., 2021).⁵ Our work is also related to a series of papers that examine the impact of government interventions during the COVID epidemic on consumption (Hoseini and Beck, 2020), though in our case we show that Indian government support, though widespread, was modest as a percentage of household expenses. One feature that distinguishes our work is that we have access to longitudinal data on households so we are able to control for variation across households. Work that is similar in purpose and used the same data as the present paper is Bertrand et al. (2020) and Bertrand et al. (2020), which examine the same data and yields similar descriptive results, and Deshpande (2020), which uses the same data to examine the impact of COVID on hours of work. Other papers on the impact of COVID on Indian income include Dhingra and Machin (2020), Lee et al. (2020), Pinto et al. (2020). These papers on India do not evaluate labor churn or decompose consumption shifts into those due to prices and to income. A paper that is similar in scope to the present paper but examined the effect of the currency crisis of 1997 on Indonesia is Frankenberg et al. (2003).

The remainder of this paper adheres to the following outline. Section 2 describes the CMIE Consumer Pyramids data. Section 3 characterizes the timing and distribution of changes in income. Our focus is on how workers tried to smooth income by switching occupations. Finally section 4 characterizes the timing and distribution of changes in

⁵There are a series of papers that try to measure the impact of COVID on economic activity using indirect measures such as night lights (Beyer et al., 2020) and NO₂ emissions (Deb et al., 2020).

consumption. We use Engel curves to decomposes the consumption response to COVID into changes due to income and prices.

2 Data

2.1 Consumer Pyramids Household Survey

Our primary data source is the Center for Monitoring Indian Economy’s Consumer Pyramids Household Survey (CPHS). This is a panel of roughly 174,000 Indian households (1.2 million members) surveyed every four months since January 2014. Households are intended to be representative at the level of urban and rural areas of regions, defined as clusters of similar districts within a state. Sampling is staggered so that roughly 25% of all households are sampled each month.

Sampling and weights. The CPHS covers nearly all states in India except for a few states in the northeast (e.g., Nagaland, Mizoram, Arunchal Pradesh) that are difficult to sample because of instability (Vyas, 2020d). CPHS divides each state into homogeneous regions, clusters of districts with similar features.⁶ Each region is divided into rural and urban strata, where rural regions are villages as defined by the Indian Census. The urban strata is further subdivided into four sub-strata defined by town size. The primary sampling units are villages and towns. Thirty villages were randomly selected from rural strata of each region. For urban strata, a random subsample of towns in each sub-strata are selected. The ultimate sampling units are households. In each selected village 16 households were selected by systematic random sampling (every n^{th} household on a street, where n is a random number between 5 and 15). In cities, 21 Census Enumeration Blocks (CEB) were randomly selected. In each CEB, 16 households were selected via systematic random sampling.

Churn and non-response. The sample churns somewhat over time Vyas (2020c). On average 2.1% of households are lost in each four month wave and 2.4% are added in each wave to replace lost households and to grow the sample over time. Weights are included to ensure that the sample remains representative of its region. Prior to the COVID epidemic, response rates were roughly 84%. Non-response was due more to inability to reach all households in the allotted 4 months for each sampling wave more than refusal to be surveyed. Separate weights are included in an attempt to correct for non-response.

During the pandemic, response rates fell (Vyas, 2020a). After states and then the central government declared a lockdown in the 3rd week of March, the survey switched temporarily from in-person to telephonic. Response rates during the 19th wave (January–April 2020) fell to just 64.4% (as low as 30% in April). The 20th wave (May–August 2020) had a response rate of just 40% (Table 1).

Nevertheless, CMIE was able to maintain the distribution of surveys across two dimensions the same as prior to the pandemic. First, the ratio of rural to urban households was roughly 35:65 pre- and post-lockdown, with only a two week deviation to 43:57 when lockdown was declared. The distribution of households across states shifted a bit in favor

⁶The features are similar agro-climactic conditions, urbanization levels, female literacy, and average household size.

Table 1: CPHS Response Rates through Lockdown

Month	Sample (numbers)	Response as % of sample
January 2020	43,352	83.3
February 2020	42,816	86.5
March 2020	43,931	58.9
April 2020	44,306	30.0
May 2020	43,352	35.1
June 2020	42,816	39.7
July 2020	43,931	44.3

Note. Data are from Vyas (2020c).

of rural states, but was roughly the same as pre-pandemic. Second, the distribution across income also remained roughly the same. The fraction earning between ₹150,000–300,000 per annum was 45% before and after lockdown. However, sampling at the extremes of the distribution did change. The share earning \geq ₹500,000 fell from 12.9% to 9%; the share earning \leq ₹150,000 increased from 23% to 29.1%. Of course, some of this change may be a reduction in income due to COVID⁷

Survey content. The survey is conducted at the household level but measures both individual and household level variables. It measures employment status, time use, and occupation for each member of the household once every four months. It measures income for each household member and the overall household and consumption for the household every month by asking members to recall income and consumption each of the last four months. The survey is conducted on a smart device and gathers data on up to 12 members of each household.

The CPHS has analogues in Indian NSSO surveys on labor statistics and on consumer expenditures. It is difficult to compare the CPHS to the consumer expenditure survey because the NSSO’s 2017–2018 consumer expenditure survey was rejected by the government and thus not released (Vyas, 2020b). The previous one was from 2011–2012. However, Abraham and Shrivastava (2019) show that the CPHS and NSSO produce similar results for male workers.

2.2 RBI inflation data

We obtain data on overall and constituents⁸ of inflation from the Reserve Bank of India (Reserve Bank of India, 2020). The data are available at the monthly level for rural and urban areas of each state; the base year is 2012. We obtain relative prices from constituents

⁷CMIE also reports a shift in the fraction of households across occupational groups. Again, this could be churning due to COVID as opposed to sampling. In any case, we address the problem by looking at within household change in occupation when analyzing labor supply.

⁸Food, clothing, fuel and light, transport and communication, intoxicants, housing, household goods and services, health and education, and others.

price indices using the rural and urban constituent weights reported by the RBI (Bhoi et al., 2020).⁹

2.3 COVID cases, lockdown rules and mobility

We obtain data on COVID cases and deaths from www.covid19India.org, which compiles reports from government sources across the country. We obtain data on national lockdown severity from the Oxford Covid-19 Government Response Tracker (Hale et al., 2020) and district-level mobility from Google Mobility Reports (Google LLC, 2021).

3 Income

3.1 Decline in income

Income. Per capita income fell sharply during the pandemic (Figure 2).¹⁰ At its low point, mean income was 40% and median income was 65% lower than average 2019 income, reflecting an enormous unanticipated shock to incomes.¹¹ Interestingly, income fell up to 20% after India reported its first case (in Kerala on 30 January 2020) but before Indian governments instituted lockdowns in the third week of March. This suggests some of the income decline was due to voluntary contraction of economic activity before mandatory business closures.¹²

Income recovered rapidly, but remains about 10% lower than 2019 levels as of October 2020, the latest data available for this paper. The recovery began in May and was V-shaped. Indeed, the recovery seem to have occurred even before the epidemic peaked (Figure 1), though it is possible the recovery drove the epidemic.

Employment rate and hours. Employment and hours conditional on employment fell sharply (Figure 3). The effects were more severe for males (20 percentage points) than females (3 percentage points), though females had much lower employment rates (< 10%) in 2019 to begin with. As with overall income, employment and wages recovered rapidly, even before the ostensible peak of the epidemic.

3.2 Distribution of decline in income

We explore the distribution of changes in income across two dimensions: occupation and household income. First, we compare changes in four high-level occupational groups: salaries employees, business persons, farm and agricultural laborers, and small traders


⁹CPI for items other than constituent g ($\neg g$) calculated as $CPI_{\neg g} = (\sum_{k \neq g} w_k CPI_k) / (\sum_{k \neq g} w_k)$, where k indexes constituents.

¹⁰We calculate per capita income as household income divided by household size. We do this to account for household level income, including from household-level businesses.

¹¹Our benchmark is the average monthly income across months. For median (mean) income, we use the median (mean) across households of average monthly income as a benchmark.


¹²To be sure, India did impose increasing international travel bans prior to the domestic lockdowns. However, is hard to imagine those were solely responsible for restrictions on domestic activity to the extent observed.

Figure 2: Changes in Income and Consumption per Capita



Note. The figure was constructed by first dividing the household income by the household size to calculate per capita income, then dividing by the state \times urban status specific mean or median 2019 income, and finally calculating monthly means or medians using individual member weights. A similar process was followed for consumption. Dashed vertical lines in January 2020, March 2020 and June 2020 indicate the month of first case (blue), the month the national lockdown started (red) and the month the national lockdown ended (green). All values are inflation adjusted in 2012 ₹.


Figure 3: Employment by Gender




Note. Includes all individuals aged 15 or above. Share employed is the number of people who report being employed relative to the labor force, retirees, and students. Daily hours worked is average hours worked conditional on working any positive hours. Dashed vertical lines in January 2020, March 2020 and June 2020 indicate the month of first case (blue), the month the national lockdown started (red) and the month the national lockdown ended (green).

Figure 4: Change in Income by Occupation categories


Panel A: Major Occupation Categories



Panel B: Constituents of Salaried Workers



(A) Panel C: Constituents of Agricultural Workers



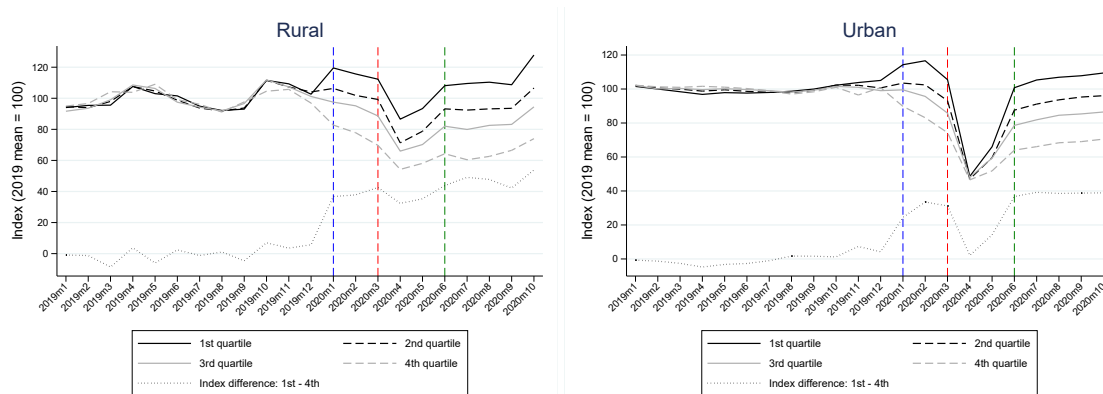
Note. An index of mean per capita income was generated by constructed by dividing the household income by the household size to calculate per capita income, then calculating the mean across India, using individual member weights, by current occupational category, and finally dividing by mean per capita income in 2019. Dashed vertical lines in January 2020, March 2020 and June 2020 indicate the month of first case (blue), the month the national lockdown started (red) and the month the national lockdown ended (green).

and daily workers. We do not disaggregate by community type (urban or rural) in our occupational breakdown. Second, we compare change in four quartiles of income based on 2019 average household income. For income, we do disaggregate by location, as urban incomes are higher. Our units for measuring income are percent of average 2019 income to abstract from variation in income levels across locations.

Income. While all occupations saw sharp drops in income, there were large differences across occupations (Figure 4). While salaried employees saw mean income fall 40% in April 2020, small traders and daily laborers (henceforth daily laborers) saw a much larger decline of 75%. Even these declines mask substantial inequality in income changes. Among salaried workers (Panel B, Figure 4), white collar workers saw declines of roughly 20%, while blue and pink collar workers saw declines of 50%. Among agricultural workers (Panel C, 4), laborers (who work on land owned by others) saw a decline almost triple that of small farmers (who own small plots of land). Organized farmers, who own larger plots, saw incomes fall in August 2020 by 40% of their average 2019 income. However, this group—and to a lesser extent small farmers—is the residual claimant on agricultural revenues, which fluctuate dramatically over even a normal year like 2019. That means the decline relative to their August 2019 baseline is closer to 10%.

Together these changes suggest that salaried white collar workers were hit far less during the COVID pandemic; while daily laborers, including in agriculture, were hit far harder. An important topic for future research is whether salaried occupations were protected because they employ high human capital workers, permitted greater remote work, or offered greater contract security.

Figure 5: Change in Income by Household Income Quartiles



Note. Individuals were assigned to income quartiles calculated using average 2019 per capita incomes. The figure was constructed by first dividing the household income by the household size to calculate per capita income, then calculating the mean within income quartiles in their state \times urban status locations, using individual member weights, and finally dividing by mean per capita income in 2019 to create an index. Dashed vertical lines in January 2020, March 2020 and June 2020 indicate the month of first case (blue), the month the national lockdown started (red) and the month the national lockdown ended (green). The line at the bottom indicates the interquartile range of income, measuring inequality of income.


While the income shock from Covid shows large, regressive inequalities across occupations, the shock is either neutral or progressive across household income. Figure 5 shows that incomes of individuals from the highest (4th) income quartile (based on 2019 average household income) fell earlier. Moreover, in rural areas the highest quartile fell

more at the peak of the economic crisis. Finally, incomes of the highest quartile have not returned to pre-crisis levels, while those in the lowest quartile have. Although we do not know why the shock was progressive, this finding is consistent with the idea that some members of higher income groups were more cautious about returning to work due to fear of catching COVID.


Employment rate and hours. Unpacking the income shock across the occupation distribution (Figure 6), we see that daily laborers and those from higher-income households saw greater increases in unemployment. Salaried workers saw the greatest decline in hours (even when employed), which is consistent with higher income workers having a lagging recovery from the initial economic crisis. After salaried workers, daily laborers saw the biggest decline in employment. A surprising finding is that the decline in hours was similar across income quartiles. Again, disaggregating employment by occupation yields different conclusions about inequality than disaggregating by income.

Figure 6: Male Employment by income Quartile and Occupational Category


Panel A: Employment rate, by occupation




Panel B: Hours, by occupation



Panel C: Employment rate, by income quartile



Panel D: Hours, by quartile



Note. Includes male individuals aged 15 or above. The occupation categories are as of 4 months prior. Share employed is the number of people who report being employed relative to the labor force, retirees, and students. Daily hours worked are medians conditional on working any positive hours. Dashed vertical lines in January 2020, March 2020 and June 2020 indicate the month of first case (blue), the month the national lockdown started (red) and the month the national lockdown ended (green).

Wages. Combining income and hours data, we are able indirectly to measure hourly wages by occupation. We do this by estimating the following regression of individual (i) income (I_{iokt}) on hours (h_{iokt}) in an occupation o separately for each location k and time


period t :

$$I_{iokt} = w_{okt}h_{iokt} + e_{iokt} \quad (1)$$

where location is defined community type (urban or rural).¹³ The coefficient on hours yield implied hourly wages.

The resulting time series (Figure 7) reveals daily workers were hit more than other occupations not just on the hours margin, but also the hourly wage margin. Their wages fell roughly 40% in rural areas and 80% in urban areas at the height of the downturn. Salaried workers, unsurprisingly see the least decline in implied hourly wage. Indeed, in rural areas, their salaries stayed constant but they were asked not to come into work enough that their implied hourly wage increased 30% during the first full month of lockdown. By October 2020, wages for non-agricultural workers were just 10% below 2019 levels. Agricultural wages still lagged by 20–30%.

Figure 7: Wages by Occupational Category



Note. Includes all males and females aged 15 or above. Wages are coefficients from a regression of wage income (not per capita household income) on hours, run separately by occupation and location. We report mean implied wages across locations, by occupation. Wages are reported conditional on positive hours. Dashed vertical lines in January 2020, March 2020 and June 2020 indicate the month of first case (blue), the month the national lockdown started (red) and the month the national lockdown ended (green).

3.3 Income smoothing via labor supply

COVID was, at least initially, an aggregate supply shock (Guerrieri et al., 2020). Individuals may have attempted to smooth income through the usual tools in their toolkit. However, many of those efforts may have been blunted by the combination of the surprise and aggregate nature of the shock and the social distancing orders. For example, it might be difficult to shift towards contractual protections (Bardhan, 1983) or modify work inputs (Binswanger and Rosenzweig, 1993) because the pandemic was unanticipated. Moreover, attempts to move to find work was impeded by the aggregate nature of the shock and government restrictions on travel.

¹³In CPHS data, individuals can only report their primary occupation. (We implement this as a regression where hours are interacted with time, occupation and location fixed effects, with no constant or main effects.) We assume the hours of work they report are entirely in that occupation.

That is not to say that there was no scope for adjustment. It is possible that individuals sought more hours at lower pay. It is possible that they shifted to other sectors or occupations in search of work. Given the aggregate nature of the shock, however, these efforts were likely to have spillover effects on other workers. We explore these adjustments and possible spillovers in this section. We start at the intensive margin.

3.3.1 Intensive margin

We explore changes at the intensive margin by examining how workers changed the number of hours they worked during COVID. It is possible that workers tried to offset the decline in wages illustrated in Figure 7 by increasing hours worked. This would be observed if the income effect was greater than the substitution effect on leisure over the observed, COVID-period wage range. Others have observed such income-smoothing adjustments in emerging economies (Moser, 1998).

To explore changes in the intensive margin of labor supplied, we report transition matrices for hours worked. First, in Table 3 we segment hours worked into bins (not employed, employed but 0 hours, (0,4] hours, (4-6] hours, (6-8] hours, and > 8 hours worked). Second, for each of two months (September 2019 and December 2019), we calculate the transition matrix of probabilities of going from each of the hours bins that month to each hours bin 4 months later (January 2020 and April 2020, respectively).¹⁴ We assign each transition matrix to the destination month, e.g., we call the December 2019–April 2020 transition matrix simply the April 2020 transition matrix. To highlight the change in transition probabilities from pre-COVID to during COVID, we subtract the monthly transition matrix for January 2020 from the matrix for April 2020. While January 2020 is supposed to be the pre-COVID transition matrix, we note there may be some complications associated with seasonality in hours worked. To mitigate seasonality effects, we leave out agricultural workers from our analysis. (We cannot simply subtract the transition matrix from 1 year prior because CPHS reports hours starting only in September 2020.) The resulting change-in-transition matrix is reported in Table 3.

We find little evidence that individuals were able to increase their hours to offset the aggregate supply shock. The first panel looks at all occupations other than agriculture to reduce the impact of seasonality on our analysis of hours. We see massive increases in unemployment in April relative to regular churn. If we add people who claimed to be employed but happen to show zero hours worked, there is a roughly 65% increase in non-work. Even among those who work, we only see declines in hours, regardless of how many hours one previously worked. These changes also suggest much of the shock was felt at the extensive rather than intensive margin, even if we put aside smoothing at the intensive margin.

When we examine each occupation separately, focusing only on workers who were in the occupation in the prior period, farms and farm workers seem to stand out. They were less likely to end up without employment or hours. They are also more likely to retain their pre-COVID hours during the pandemic. One should take these findings with a grain of salt: there is a great deal of seasonal variation in agricultural work, and April is typi-

¹⁴We can only focus on two months because CPHS does not have hours data prior to September 2019.

cally a high mark for income prior to the pandemic (see Figure 7). We omit agricultural workers in the first panel for this reason.

Table 2: Change in hourly transition matrices for April 2020, by occupation.

	Not employed	0 hours	(0-4] hours	(4-6] hours	(6-8] hours	>8 hours
All but agricultural workers						
0 hours	42.43	41.50	0.00	-6.64	-64.50	-12.79
(0-4] hours	61.20	7.59	-29.14	-5.47	-26.74	-7.44
(4-6] hours	46.28	20.88	-1.22	-12.07	-39.03	-14.84
(6-8] hours	47.71	27.07	2.43	1.18	-58.58	-19.81
>8 hours	36.39	39.29	3.63	4.16	-41.40	-42.07
Salaried employees						
0 hours	20.61	40.33	0.00	-12.79	-41.51	-6.64
(0-4] hours	76.45	7.08	-3.61	-16.64	-36.38	-26.91
(4-6] hours	32.79	22.42	0.21	-16.24	-34.30	-4.88
(6-8] hours	29.48	34.59	4.10	3.44	-58.32	-13.29
>8 hours	29.34	42.74	2.80	5.30	-49.73	-30.45
Farm & farm workers						
0 hours	7.90	-33.56	0.00	43.05	0.72	-18.10
(0-4] hours	38.79	7.65	11.00	-16.39	-35.88	-5.17
(4-6] hours	29.20	6.41	17.46	5.76	-39.46	-19.37
(6-8] hours	35.55	10.84	7.98	9.99	-47.20	-17.17
>8 hours	20.35	9.50	8.87	23.41	-16.97	-45.17
Business workers						
0 hours	30.41	36.52	0.00	-3.31	-32.66	-30.95
(0-4] hours	57.56	8.08	-25.34	-7.94	-27.72	-4.65
(4-6] hours	43.77	17.01	1.81	12.24	-46.48	-28.35
(6-8] hours	36.46	29.63	6.86	4.46	-49.77	-27.65
>8 hours	25.63	36.69	8.77	9.48	-34.18	-46.38
Small traders & daily laborers						
0 hours	100.00	0.00	0.00	-6.27	-73.54	-20.19
(0-4] hours	59.43	9.40	-34.83	-3.08	-25.36	-5.56
(4-6] hours	42.12	15.34	-1.55	-8.15	-31.97	-15.79
(6-8] hours	54.00	13.48	5.27	0.78	-53.90	-19.63
>8 hours	47.05	27.64	1.93	6.80	-39.48	-43.94

Note. This table is generated as follows. First, we segment hours worked into bins (not employed, employed but 0 hours, (0,4] hours, (4-6] hours, (6-8] hours, and > 8 hours worked). Second, for each of two months (September 2019 and December 2019), we calculate the transition matrix of probabilities of going from each of the hours bins that month to each hours bin 4 months later (January 2020 and April 2020, respectively). We assign this transition matrix to the destination month, e.g., we call the December 2019 - April 2020 transition matrix the April 2020 transition matrix. These transition matrices are limited to workers working in the occupation(s) indicated in the title of each panel in the origin or destination month. Third, we difference the monthly transition matrices, i.e., we subtract the transition matrix for January 2020 from the transition matrix for April 2020.

3.3.2 Extensive margin - unemployment and occupational churn

With respect to adjustments at the margin of employment, we look at two behaviors: whether a person works and what sector he or she works in. To capture both we define

five exhaustive states: not employed¹⁵ now and in the last period,¹⁶ not employed now but employed last period, employed in same occupational category as the last period, employed in a different occupational category in the last period, employed but unemployed or OLF in the last period. We further subdivided (the previously unemployed → now employed) and the (previously employed → now employed in different occupation) into two categories: employed in the same sector one was previously in (switchback) and employed in a new sector (no switchback). This extra partition is intended to determine possible loss of productivity due to adaptation. Presumably, individuals who switch to an occupation that is new are likely to have lower productivity either relative to an occupation in which they previously worked in.

Figure 8 shows that there was an increase in unemployment and exit from the labor force during the pandemic. In addition, there appears to be an increase in switching occupations or occupational churn. When the economy resumed activity, many people switched back to occupations in which they previously worked, whether they came from the unemployed state or another occupation. However, there does appear to be a higher fraction of people working in a new occupation relative to, say, 2018 or 2019. It remains to be explored whether this occupation shifting will significantly affect longer-term labor productivity.

Occupational churn may be asymmetric, e.g., it may be possible to go from being salaried to being a daily worker, but not vice versa. A simple model that might explain this is a Roy model of sorting across sectors based on skill and skill prices. If skills are positively correlated and skill prices differ across sections, then it is possible that there is initial sorting into jobs based on comparative advantage. When there is a labor demand shock and skills are constant over time, however, it is possible workers shift to other sectors in which they previously had an absolute advantage, but not a comparative advantage. During the pandemic, skill price changes might convert those absolute advantages into comparative advantages. While this is not the only model that explains transitions, it is helpful to keep in mind as we explore transitions.

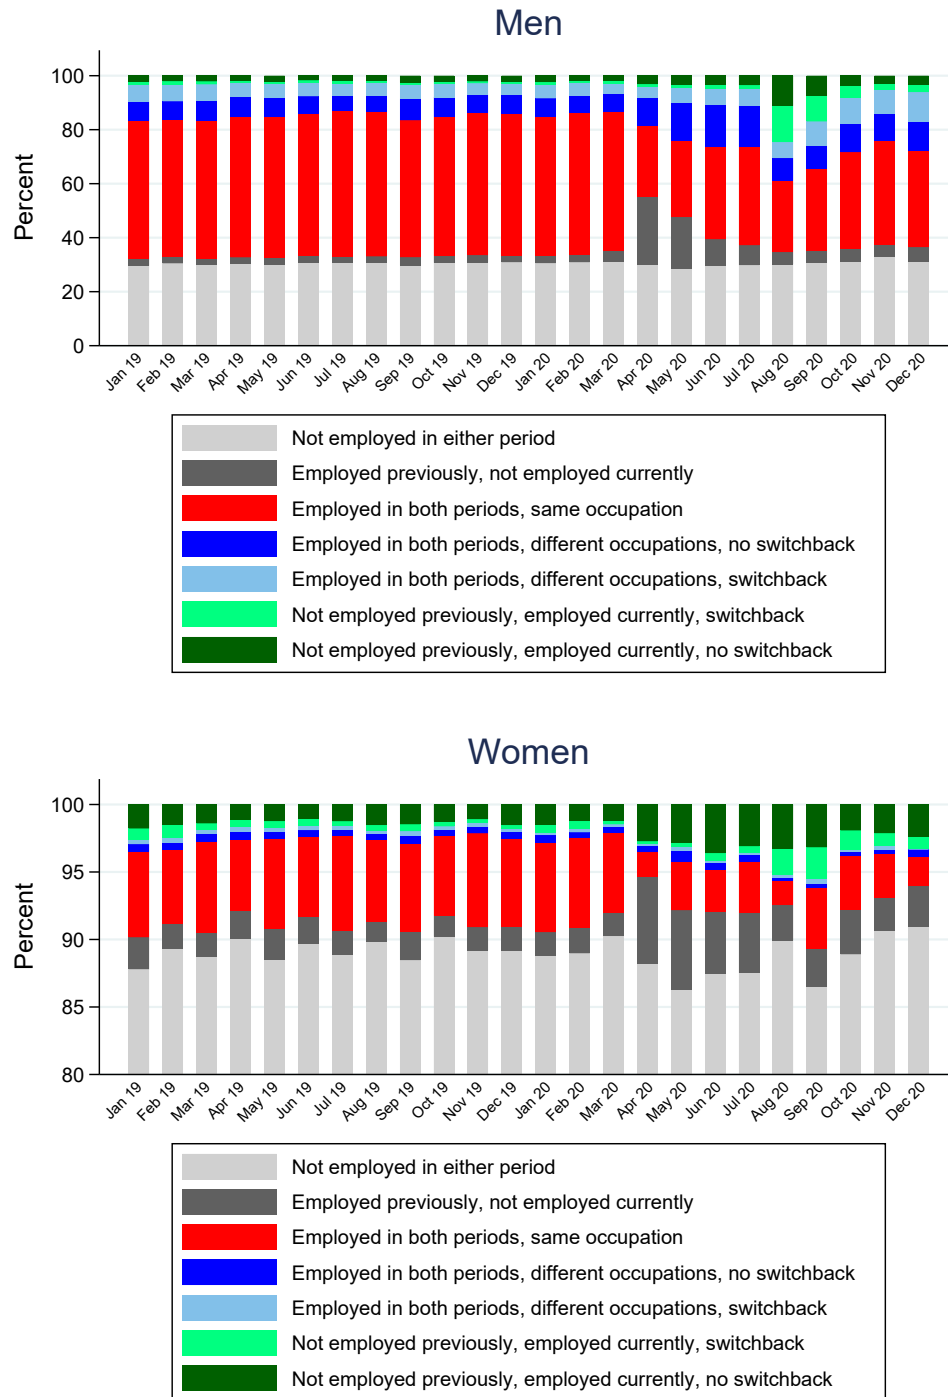
To explore asymmetric changes, we switch from reporting stacked bar charts to transition matrices in Table 3. We define five states: not employed, and employed in each of the four occupations we reported before (salaried employee, farmer and farm laborer, business person, and small trader and daily wage earner). For each month, we can calculate the transition matrix of probabilities of going from each of the states four months ago to each state the selected month. To better highlight the change in transition probabilities, we difference monthly transition matrices across year, e.g., the change in April 2020 transitions will show the difference between the April 2020 and April 2019 transition matrices. To simplify our exposition, we group months into quarters and report the change in transitions for each quarter of each year.

The results reveal four things about the downturn at the start of the epidemic. First, there was greater transition to unemployment among individuals who are in lowest income occupation, i.e., daily laborers. Second, agriculture was least disrupted; while un-

¹⁵We define not employed as out of the labor force (OLF) or unemployed, categories found in the CPHS data set.

¹⁶The last period is defined as 4 months ago, which is the last time the member was surveyed in the CPHS.

Figure 8: Labor force status over time



Note. These figures were constructed by, first, categorizing each member of each household into five states in each month they are observed: not employed now and in the last period, not employed now but employed last period, employed in same occupational category as the last period, employed in a different occupational category in the last period, employed but unemployed or OLF in the last period. (We define not employed as out of the labor force (OLF) or unemployed, categories found in the CPHS data set. The last period is defined as 4 months ago, which is the last time the member was surveyed in the CPHS.) We then calculate the fraction of the observed members in each state in each month. Each figure includes only those members aged 18–65. Switchbacks are measured by examining whether the individuals switches to a sector they had previously worked in either four months or one year previously. Note that the female graph starts at 80%. Dashed vertical lines in January 2020, March 2020 and June 2020 indicate the month of first case (blue), the month the national lockdown started (red) and the month the national lockdown ended (green).

employment increases, there was a smaller outflow from agriculture than other occupations. Unsurprisingly, it is also the occupation that saw the most inflow of workers. This is not merely a seasonal effect because we subtract the transition matrix from Q2 2019. This finding is consistent with reports of a massive migration from cities to rural areas when the country began to loosen the lockdown in May Malani et al. (2020). Third, salaried persons and business persons seemed to have relied on occupational transition to protect consumption more than daily laborers did. Fourth, the recovery in Q3 2020 largely followed these patterns, but dramatically scaled back transitions to unemployment, especially for daily laborers.

Table 3: Change in Occupational Transition Matrices for Q2 and Q3 2020

	Salaried employees	Business persons	Farmers & farm workers	Small traders & daily workers	Not employed
Q2 2020					
Salaried employees	-36.16	3.84	7.51	1.99	22.81
Business persons	6.58	-32.87	6.63	-0.94	20.61
Farmers & farm laborers	1.93	2.33	-21.31	-2.15	19.20
Small traders & daily wage earners	3.37	1.51	8.75	-49.67	36.04
Q3 2020					
Salaried employees	-33.54	7.11	8.15	6.87	11.41
Business persons	6.89	-26.35	6.73	6.10	6.64
Farmers & farm laborers	1.99	3.48	-16.74	3.89	7.38
Small traders & daily wage earners	2.58	4.38	8.39	-24.86	9.51

Notes. Q2 is April–June; Q3 is July–September. We define five states: not employed, employed in each of the 4 occupations we reported before (salaried employee, farmer and farm laborer, business person, and small trader and daily wage earner). For each month, we calculate the transition matrix of probabilities of going from each of the states in the chosen month to each state after four months as the observed fraction of the population that made that transition. The fraction is calculated only among household members between ages 18–65; moreover, these members are weighted to represent the national average. We assign this transition matrix to the destination month, e.g., we call the December 2019–April 2020 transition matrix the April 2020 transition matrix. To obtain the change in transition probabilities, we difference monthly transition matrices across years, e.g., the April 2020 change-in-transition matrix is the difference between the April 2020 and April 2019 transition matrices. Quarterly change in transition matrices are the equally weighted change-in-transition matrices for months that compose the quarter.

Occupational churn occurred even before the pandemic. However, the net benefit of occupational churn can change in the presence of aggregate shock like COVID. For example, with a shock in just one occupation or sector, an individual in a origin occupation that suffers a negative shock can move to another destination occupation that has not suffered the same shock. While the additional labor supply in that second occupation will reduce wages there, the effect may be minor. With a shock to all occupations, however, this destination supply effect can be magnified if the displacement is asymmetric, i.e., some destination occupations see more entry of labor. This can negatively affect the private return from occupational switching to protect consumption, as well as the external effects of churn on other workers in the destination occupation.

To address this challenge, we calculate the private benefit of occupational churn during an aggregate shock like COVID. Define t as the pre-COVID period, $t + 1$ as the COVID period, j as the pre-COVID occupation, and $j' \neq j$ as the COVID period occupation for

individuals who switched occupations. Income pre-COVID is $Y_{ijk,t+1} = w_{jk,t} h_{ijk,t}$, where k indexes location. Income for individuals who switched occupations during COVID is $Y_{ij'k,t+1} = w_{j'k,t+1} h_{ij'k,t+1}$. Consider a person who has switched occupations during COVID. We want to calculate how much additional income she made by switching occupations and how that changed during the pandemic.

These calculations require specifying counterfactuals. For example, if the person were a typical person in her prior occupation per-COVID, she may have made what the typical person in that old occupation makes during COVID. But the fact that the person switched suggest she is probably worse than the typical person remaining in that prior sector. Conversely, if one's model for transitions is the changing comparative advantage model described above, then it is also possible that the transitioning person is actually making a higher daily income than the typical person in the destination district because they are greater skill than others in the destination occupation. These problems are worse when one considers that switchers may not be representative of either district on hours as well as hourly wage.

To calculate the counterfactual against which to value churn, we employ the following model of wages in sector j at time t :

$$Y_{ijkt} = \beta_{jt} X_{ijkt} + \gamma_{jt} I_{ik,t}(j) + e_{ijkt} \quad (2)$$

where $I_{ik,t}(j)$ is an indicator for whether the person just switched into sector j . The coefficient γ_{jt} captures the lower wage that switchers may earn because they lack experience in the sector. Our goal will be to estimate the value of switching from sector j to j' :

$$V_{i,j,j',kt} = \hat{\beta}_{j'k} X_{ij'kt} + \hat{\gamma}_{j'kt} - \hat{\beta}_{jk} X_{ijkt} \quad (3)$$

Because switching occurred with great frequency even before COVID, we want to estimate the wage equation and the value of switching with cross-sectional regression in 4 periods: (i) Q2 of 2018 and 2019, (ii) Q3 of 2018 and 2019, (iii) Q2 of 2020, and (iv) Q3 of 2020. Comparing the Q2 (Q3) 2020 gain to the Q2 (Q3) 2018/2019 gain tells us the change in benefits to switching during the downturn (recover) period of the pandemic.

The first problem we face is that there is selection into sectors. To address this, we use a Roy model framework French and Taber (2011); allow workers a choice of sectors to enter, a decision captured via a multinomial probit (e.g., Table 4 for Q2 in 2018/19 and Q2 2020); and use experience in each sector as our exclusion restriction for each sector. The second problem is that we do not observe reservation wages (i.e., implicit wages in the unemployed sector). To capture these we estimate a version of the model with just two sectors, employed and unemployed, and back out estimates of reservation wages using a probit first stage and a single wage equation for the employed sector. The result is that the gain from switching to unemployment is the same regardless of which sector a person switches from (Table 5).

Table 4: First Stage Probit Estimation for Roy Model

	Q2 2018/19			Q2 2020				
	Unemployed (1)	Agriculture (2)	Services (3)	Industry (4)	Unemployed (5)	Agriculture (6)	Services (7)	Industry (8)
Age	10.393*** (0.292)	303.937*** (10.014)	485.108*** (14.919)	628.563*** (38.356)	14.051*** (1.483)	332.603*** (19.904)	477.615*** (56.784)	571.864*** (52.090)
Age ²	-0.124*** (0.004)	-3.244*** (0.108)	-5.168*** (0.184)	-6.387*** (0.442)	-0.119*** (0.020)	-3.463*** (0.335)	-5.186*** (0.765)	-5.780*** (0.578)
1st Std	3.543 (8.669)	92.945 (138.437)	413.882** (191.669)	598.316* (351.494)	29.175 (40.352)	-1180.421*** (165.250)	-99.277 (511.705)	1068.922 (1951.425)
2nd Std	-15.495 (10.425)	113.967** (49.950)	598.603*** (126.335)	946.825*** (228.994)	-42.952* (23.954)	-91.814 (585.589)	304.729 (724.896)	556.072 (1152.833)
3rd Std	-26.575*** (6.734)	220.365 (251.347)	479.966*** (153.730)	951.879 (622.674)	-44.045 (29.564)	53.220 (338.575)	319.889 (595.322)	732.489 (2246.287)
4th Std	-29.706*** (4.568)	98.190 (84.812)	447.748*** (62.561)	915.325*** (212.277)	-91.522*** (25.808)	-266.191* (145.635)	128.826 (359.891)	501.000 (1056.414)
5th Std	-30.427*** (9.698)	21.511 (184.571)	988.566*** (68.998)	941.269*** (231.956)	-86.086** (36.263)	796.294*** (243.249)	440.444 (312.979)	801.485 (785.584)
6th Std	-33.526*** (10.611)	240.003* (122.794)	1346.629*** (37.802)	1753.398*** (405.368)	-17.059 (21.334)	356.102 (384.974)	29.560 (264.044)	1241.541 (1440.901)
7th Std	1.664 (14.478)	153.359 (148.095)	1415.236*** (127.699)	2055.438*** (208.649)	-50.329* (26.853)	479.465* (260.935)	405.812 (536.957)	1933.026 (1414.885)
8th Std	-4.479 (9.545)	-72.942 (159.822)	1584.407*** (53.466)	1714.363*** (216.043)	-38.860 (70.618)	773.326*** (240.960)	889.926*** (343.587)	1309.487 (1014.634)
9th Std	7.524 (8.455)	292.624 (232.185)	1738.018*** (115.171)	2505.878*** (231.946)	-40.338 (37.035)	516.000*** (200.001)	579.392** (265.167)	2264.652** (1014.984)
10th Std	18.870** (9.237)	540.657*** (143.926)	2440.334*** (99.857)	2351.614*** (254.884)	28.352 (41.930)	1123.072*** (295.501)	736.058** (322.754)	2480.915* (1362.398)
11th Std	44.712** (21.161)	373.701 (319.809)	2114.283*** (284.387)	2655.184*** (690.335)	-26.607 (51.378)	-175.630 (390.067)	572.456 (475.497)	2681.988** (1265.447)
12th Std	104.503*** (10.076)	101.949 (195.368)	3196.899*** (52.922)	3557.251*** (185.168)	33.429 (42.478)	1261.393*** (297.491)	1426.056*** (337.546)	3165.361*** (979.361)
Graduate	379.421*** (41.907)	789.554*** (288.615)	6247.938*** (174.296)	7978.626*** (226.439)	206.586*** (48.155)	949.261*** (320.313)	3121.168*** (441.801)	5185.561*** (1299.759)
Post-Graduate	688.040*** (31.373)	2414.386*** (502.920)	11375.825*** (747.915)	14911.962*** (433.038)	175.618 (166.704)	1633.740 (1334.613)	4929.191*** (794.054)	9188.775*** (1407.333)
State Mfg Wage	-0.002*** (0.001)	0.090*** (0.008)	0.738*** (0.009)	0.054*** (0.014)	0.013** (0.005)	0.251*** (0.048)	0.804*** (0.040)	0.007 (0.064)
State Svc Wage	0.003* (0.001)	0.107*** (0.010)	-0.093*** (0.008)	0.550*** (0.026)	0.006 (0.005)	0.308*** (0.060)	0.043 (0.026)	0.712*** (0.083)
Observations								
R ²								

Table 5: Gain in Occupational Shifting

	Unemployment	Agriculture	Manufacturing	Services
Q2 2018/2019				
Unemployment		3.481001	4.374509	3.890053
Agriculture	1.280828		3.321097	2.910486
Manufacturing	1.280828	3.027743		3.305666
Services	1.280828	3.659595	2.567094	
Q2 2020				
Unemployment		3.608613	4.94178	3.877368
Agriculture	0.818973		3.887736	2.894898
Manufacturing	0.818973	2.601779		2.847056
Services	0.818973	3.226935	2.831576	
Q3 2018/2019				
Unemployment		3.437133	4.310505	3.823856
Agriculture	1.302244		3.140624	2.778371
Manufacturing	1.302244	2.926753		3.321269
Services	1.302244	2.567094	3.659595	
Q3 2020				
Unemployment		3.400233	2.562705	2.346838
Agriculture	1.075048		2.742455	2.330128
Manufacturing	1.075048	2.635401		3.079273
Services	1.075048	3.206155	2.444379	

As expected, there are positive gains to switchers from switching sectors. However, the gain to switching to unemployment, for those who do, are lower than those switching to other sectors. Relatedly, the gain from switching out of unemployment into each employed sector is typically greater than the gain to entering that sector from any other employed sector.

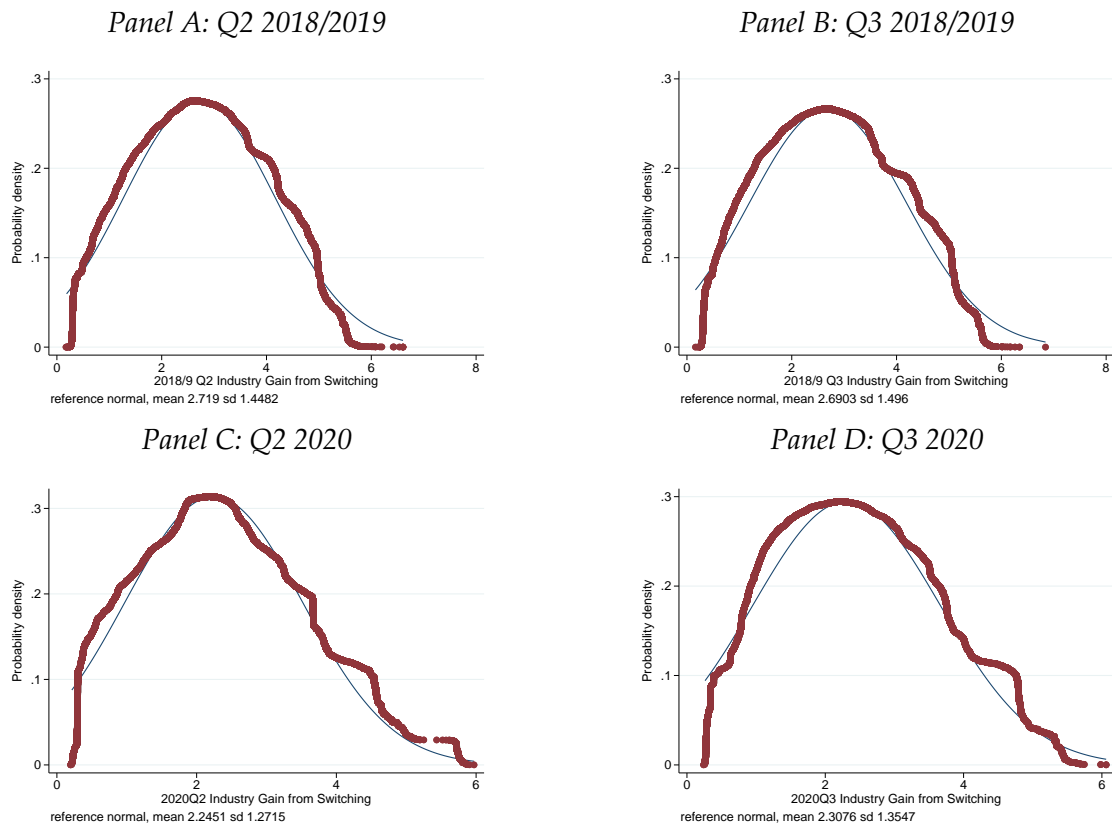
More relevant for COVID, we find that the gains to switching declined during COVID, whether one examines Q2 or Q3 of 2020 (Table 5 and Figure ??). Of particular interest is the change in reservation wage, which falls by 0.3 log points in Q2 and 0.3 log points in Q3 2020. If people were worried about going to work due to the risk of catching COVID, one would expect the reservation wage to rise, not fall. Perhaps the decline in household income caused consumption (marginal utility of consumption) in the unemployed sector to rise (fall) more than income and thus consumption in the employed sectors.

4 Consumption

4.1 Decline and distribution of decline in consumption

Indian households experienced a sharp decline in consumption at the start of the pandemic. While the average decline matched the average decline in income, the median decline in consumption was less severe than the median decline in income: median expenditure fell 40% while median income fell over 65% (Figure 2). Like income, expenditures fell before lockdowns started and recovered even before cases appear to have peaked. Consumption is now roughly 20% below the 2019 level.

Figure 9: Distribution of Gains to Switching



Note. Distribution of gains from switching is calculated over the sample of individuals who actually switched in the data. It is calculated using (3). Black line is normal approximation to maroon plot of estimated gains.

Although there are major disparities in mean income shocks across occupations (Figure 4), there was a largely uniform reduction in consumption across occupations (Figures 10), at least when the units are percent change rather than absolute change. There are roughly 20 percentage point variations in the decline in consumption across income quartiles (a) in urban areas and (b) in all areas during the recovery, but even these gaps are small relative to income gaps (Figure 11). This suggests that households were able to smooth consumption somewhat despite the fact that the pandemic was a large unanticipated, aggregate supply shock.

4.2 Consumption smoothing

We conducted basic tests of consumption smoothing, i.e., regressions of differences in consumption on differences in income, a la Cochrane (1991) and Townsend (1994). We do not expect perfect consumption smoothing: even prior to COVID, smoothing was imperfect. Instead, we test to see if households were less able to smooth during COVID. We estimate a regression of the form

$$\log c_{ikt} = \mu_i + \alpha \Delta d_{kt} + \gamma \log M_{ikt} + \pi(\log M_{ikt} \times I(2020)) + \varepsilon_{ikt} \quad (4)$$


where c_{ikt} is consumption by household i in location k , μ_i are household fixed effects, d_{kt} is a measure of the aggregate shock (proxied by location average consumption as in Townsend (1994)); M_{it} is idiosyncratic income, and $I(2020)$ is an indicator for 2020. Here γ measures risk smoothing (with $\gamma = 0$ implying full-risk sharing) and π measures whether COVID affected the ability to smooth consumption.

While we formally reject full-risk sharing, the magnitude of this rejection is somewhat small—consistent with prior literature such as Townsend (1994)—and the differential effect of COVID itself was relatively minor. Pre-COVID, a 10% income fall was associated with a 0.99% decline in consumption (Table 6). During COVID, it was associated with a 0.104% decline in consumption. Urban areas (column (4)) show less consumption smoothing than rural areas (column (3)) pre-COVID, but similarly slight increases in difficulty smoothing during COVID. To get a sense of whether the aggregate nature of the shock made a difference to smoothing, we estimate a version of this regression without the control for the aggregate shock. We find that aggregate shocks meaningfully affect smoothing, but that the aggregate shock had a larger role during COVID.¹⁷

Finally we examine how consumption changed in specific categories such as food, clothing, and education. The data show that households reduced their consumption of objects such as appliances and clothing (which fell > 90%) far more than they reduced food (fell 25%) and fuel (fell 40%). The massive decline in educational investments (fell 90%) is consistent with the finding in Chetty and Looney (2007) of inefficient smoothing because it sacrifices future income potential. Although this variation in where consumers cut back can be interpreted as Maslovian prioritization, we use the framework of Engel curves to understand these changes in the next section.


¹⁷The ratio of the coefficient on the interaction to the coefficient on \ln income in specification (1) is 0.047. It is 0.183 in specification (2), which omits a control for aggregate consumption.

Figure 10: Change in mean consumption by major occupation category



Note. The figure was constructed by first dividing the household expenditure by the household size to calculate per capita expenditure, then calculating the mean across India, using individual member weights, by occupational category, and finally dividing by mean per capita expenditure in 2019 to create an index. Dashed vertical lines in January 2020, March 2020 and June 2020 indicate the month of first case (blue), the month the national lockdown started (red) and the month the national lockdown ended (green).

Figure 11: Change in mean rural (left) and urban (right) consumption by quartiles of 2019 household income



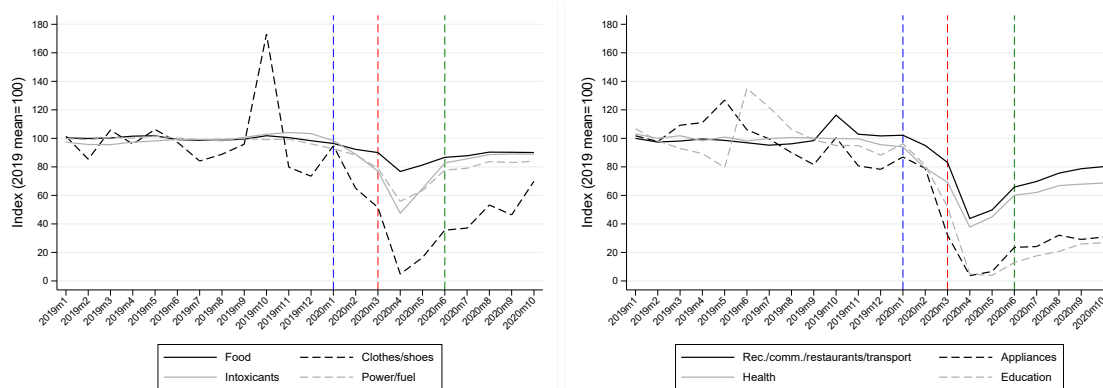
Note. Individuals were assigned to income quartiles calculated using average 2019 per capita incomes. The figure was constructed by first dividing the household expenditure by the household size to calculate per capita expenditure, then calculating the mean within income quartiles specific to their state \times urban status location, using individual member weights, and finally dividing by mean per capita expenditure in 2019 to create an index. Dashed vertical lines in January 2020, March 2020 and June 2020 indicate the month of first case (blue), the month the national lockdown started (red) and the month the national lockdown ended (green).

Table 6: Test of Consumption Smoothing

	(1)	(2)	(3) Rural	(4) Urban
Ln(ave. consumption at state \times community type level)	0.6787*** (0.0519)		0.8672*** (0.0174)	0.5612*** (0.0468)
Ln(income)	0.0944*** (0.0132)	0.1315*** (0.0155)	0.0520*** (0.0077)	0.1463*** (0.0152)
Ln(income) \times year=2020	-0.0045*** (0.0010)	-0.0241*** (0.0014)	-0.0015** (0.0006)	-0.0060*** (0.0012)
<i>N</i>	2677306	2677306	896072	1781234
<i>R</i> ²	0.728	0.701	0.720	0.730

Notes: The regressions covers years all of 2019 and January–October 2020. Household fixed effects included. Standard errors clustered at the state \times urban level. Specification (2) drops our measure of aggregate consumption, (3) only includes observations on rural households, and (4) only includes observations on urban households. Significance levels: * 10% ** 5% *** 1%.

Figure 12: Consumption Changes by Category



Note. The figure was constructed by first dividing the household expenditure in each category by the household size to calculate per capita expenditure in the category, then calculating the mean across India using individual member weights, and finally dividing by mean per capita expenditure in 2019 in that category to create indices. Dashed vertical lines in January 2020, March 2020 and June 2020 indicate the month of first case (blue), the month the national lockdown started (red) and the month the national lockdown ended (green).

4.3 Decomposing consumption changes

The causes of changes in consumption observed during COVID can crudely be broken down into changes in prices, change in income, and changes in preferences. We assume that COVID was too sudden and the economic shock too short to reflect changes in preferences. Instead, we focus on price and income as mediators for the disparate consumption changes we observed across goods.

The CPHS data allows us to observe income and consumption before and during COVID. We also have state \times month price indices for different categories of consumption from the Reserve Bank of India (RBI) price, but not during the height of the economic crisis. Specifically, we lack data on good category specific CPI from March–August 2020 and data on overall CPI for March–June 2020. In the spirit of Costa (2001) and Hamilton (2001), we assume that Engel curves can provide a structure to understand unobserved price changes.

4.3.1 Methods

Our analysis has two steps. First, we estimate Engel curves using data largely from pre-COVID in order to identify the preference parameters for each good category \times location. Second, we use our estimates of Engel curves and observed income during COVID to predict budget shares during COVID due to income changes. Then we infer that the gap between our predicted and actually observed budget shares during COVID reflect changes in the price of goods.

We start with the Working-Lesser single demand function from Deaton and Muellbauer (1980)’s Almost Ideal Demand System

$$\omega_{i,j,k,t} = \phi + \gamma(\ln P_{j,k,t} - \ln P_{j,-k,t}) + \beta(\ln Y_{i,j,t} - \ln P_{j,t}) + \theta X_{i,j,t} + u_{i,j,t} \quad (5)$$

where i indexes households, j location (homogeneous region \times urban/rural), t is time, k indicates some category of good, and $-k$ is categories other than k . $\omega_{i,j,k,t}$ is the share of a household’s budget spend in category k , $\ln P_{j,k,t} - \ln P_{j,-k,t}$ is the log of relative price of category k , $Y_{i,j,t}$ is income, and $X_{i,j,t}$ are covariates that influence preferences, including household size.

Our first step is to estimate preference parameters $(\phi, \gamma, \beta, \theta)$ by estimating the regression equation above using data from 2018, 2019 and August–October 2020. Our units are households. We control for average age of household members and household size. We employ OLS and weight each household in proportion to the number of households it represents in the nation. Because we will not be using the standard error of parameter estimates in our numerical exercise, we use robust standard errors without clustering.

Certain issues arise in running this regression. First, the categories of goods that the RBI uses does not always match the categories the CPHS uses. For example, the RBI has not released data on housing or education. Moreover, because the number of goods categories that the RBI and CPHS reports differ, we do not know if the miscellaneous category across those two sources are the same. We focus on the categories that the RBI has released. Second, expenditures sometimes exceed and other times fall short of income. This may reflect borrowing or savings, respectively. Third, there may be error in the price

indices that the government generates. Although we are motivated by a substantial literature that uses Engel curves to obtain better price indices, we sidestep this issue as our goal is to fill gaps in the CPI during the worst of the COVID downturn.

Our second step is to combine observed income and expenditures in different goods categories with our regression estimates to back out price changes that household must have experienced during COVID. Recall we do this because we lack certain CPI data during the downturn. We can derive our exact calculation by taking the total derivative of (5) and dividing by $d\omega_{i,j,k,t}$:

$$1 = \gamma_k \left(\frac{d \ln P_{k,j,t} - d \ln P_{-k,j,t}}{d\omega_{i,j,k,t}} \right) + \beta_k \frac{d \ln Y_{i,j,t}}{d\omega_{i,j,k,t}} - \beta_k \frac{d \ln P_{j,t}}{d\omega_{i,j,k,t}} + du_{i,j,k,t}. \quad (6)$$

There are two scenarios of missing-price data that we address. For some months we lack data on category-specific CPI, but have data on overall CPI. In this case, we can derive relative price changes from the equation above using:

$$\frac{\hat{\gamma}_k dE[(\ln P_{k,j,t} - \ln P_{-k,j,t})]}{d\bar{\omega}_{i,j,k,t}} = 1 - \frac{\hat{\beta}_k d \ln \bar{Y}_{i,j,t}}{d\bar{\omega}_{i,j,k,t}} - \frac{\hat{\beta}_k d \ln \bar{P}_{j,t}}{d\bar{\omega}_{i,j,k,t}} \quad (7)$$

where bars and $E[z]$ indicate sample means. Here, the right hand side is calculated and the left hand-side is inferred. The second scenario is when we have neither overall or category-specific CPI. In this case, we can only calculate:

$$\frac{\hat{\gamma}_k dE[(\ln P_{k,j,t} - \ln P_{-k,j,t})]}{d\bar{\omega}_{i,j,k,t}} + \frac{\hat{\beta}_k d \ln \bar{P}_{j,t}}{d\bar{\omega}_{i,j,k,t}} = 1 - \frac{\hat{\beta}_k d \ln \bar{Y}_{i,j,t}}{d\bar{\omega}_{i,j,k,t}}. \quad (8)$$

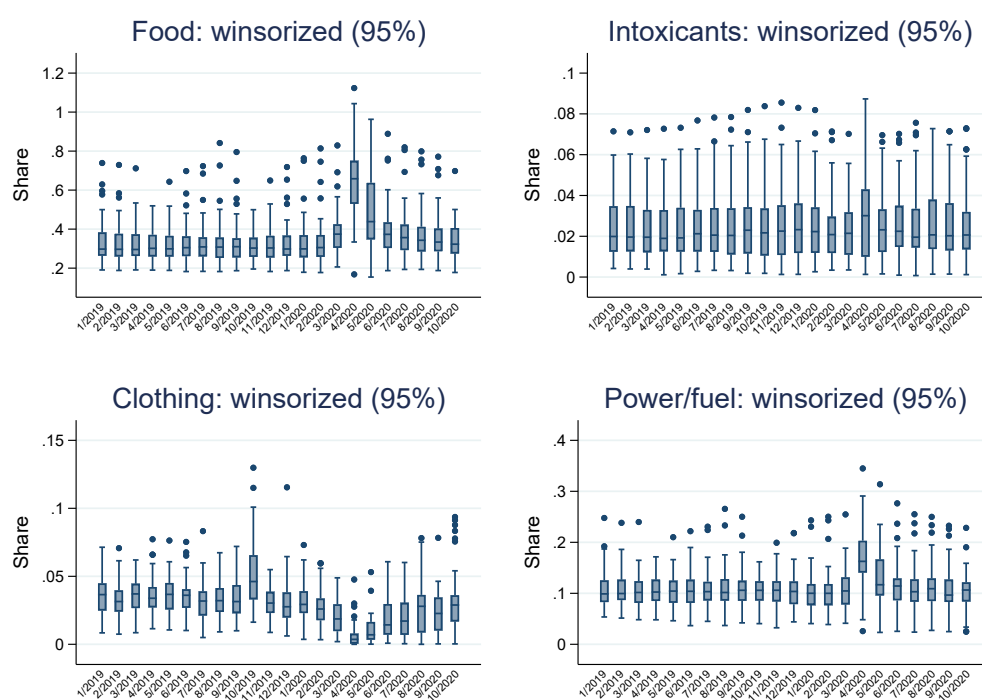
We cannot distinguish changes due to purchasing power and relative price changes.

4.3.2 Results

COVID had disparate impacts on budget shares across product categories. Figure 13 shows the distribution of budget shares over time for the four product categories for which are able to match CPI data. The food, fuel and (to a lesser extent) intoxicant share of budgets increased, while the clothing share fell. In some cases, food expenditure exceed income multiple times, suggesting households tapped savings or borrowed. However, we cannot rule out that income reporting was the problem. Because these cases were few, we winsorize the data at 95% in our analysis.

Price v. income shocks. We show the relative contribution of income and the combined effect of purchasing power and relative prices during the worse of the COVID-induced economic shock in Figure 14A. Regression coefficients are presented below each subplot. We find that income changes explain the increase in budget shares of food; price changes had little role. In short, Engels law ($\beta = -0.26$) played a major role in explaining the increase in the food share of expenditures during COVID. The smaller positive increase intoxicant share of household budgets reflect a negative shock from purchasing power or relative prices largely offsetting a positive effect from nominal income. (Both

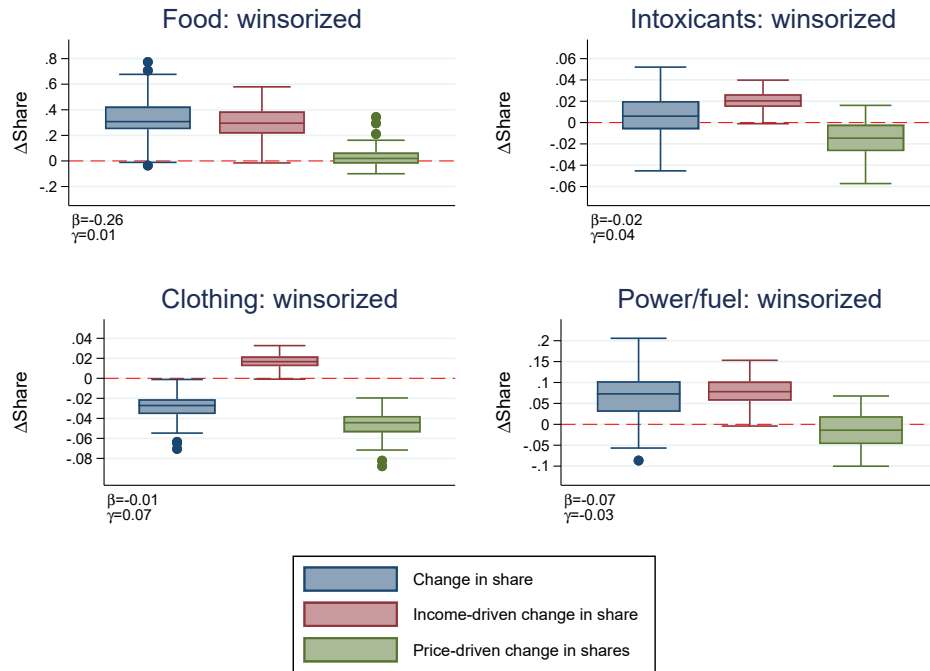
Figure 13: Budget Shares for Categories over Time



Note. Budget shares are winsorized at the 95th percentile.

Figure 14: Change in budget share and its components, by product category and date

Panel A: April 2019 to April 2020



Panel B: August 2019 to August 2020

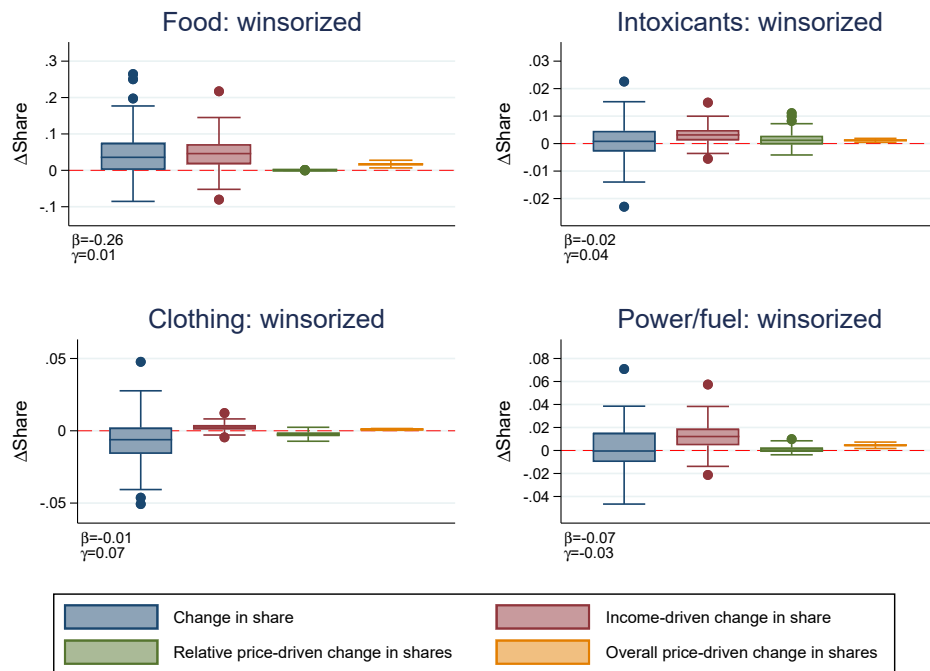
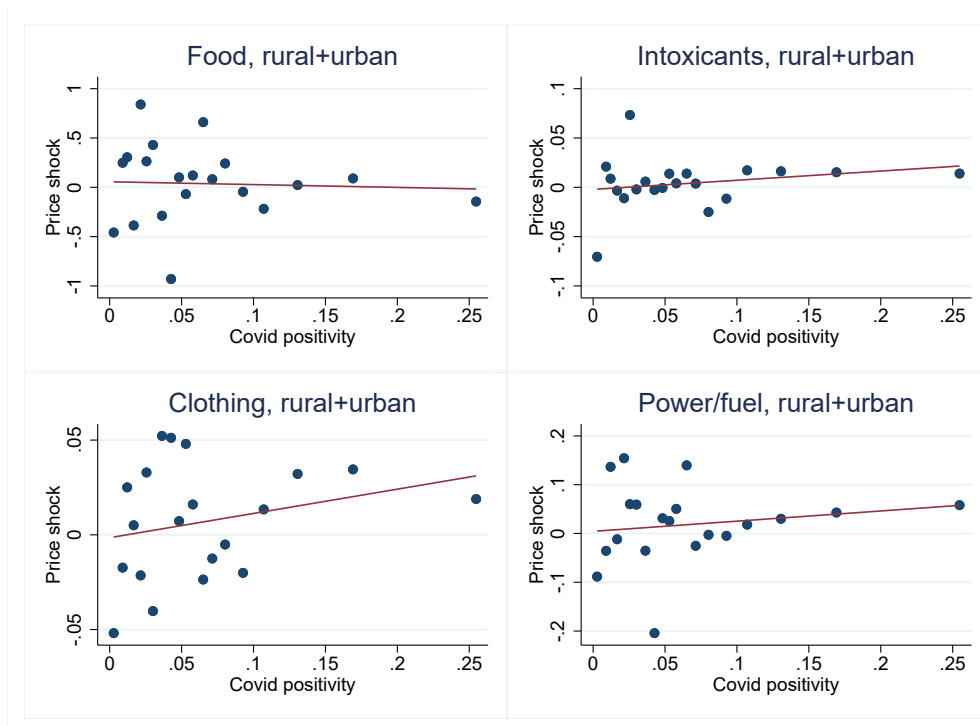
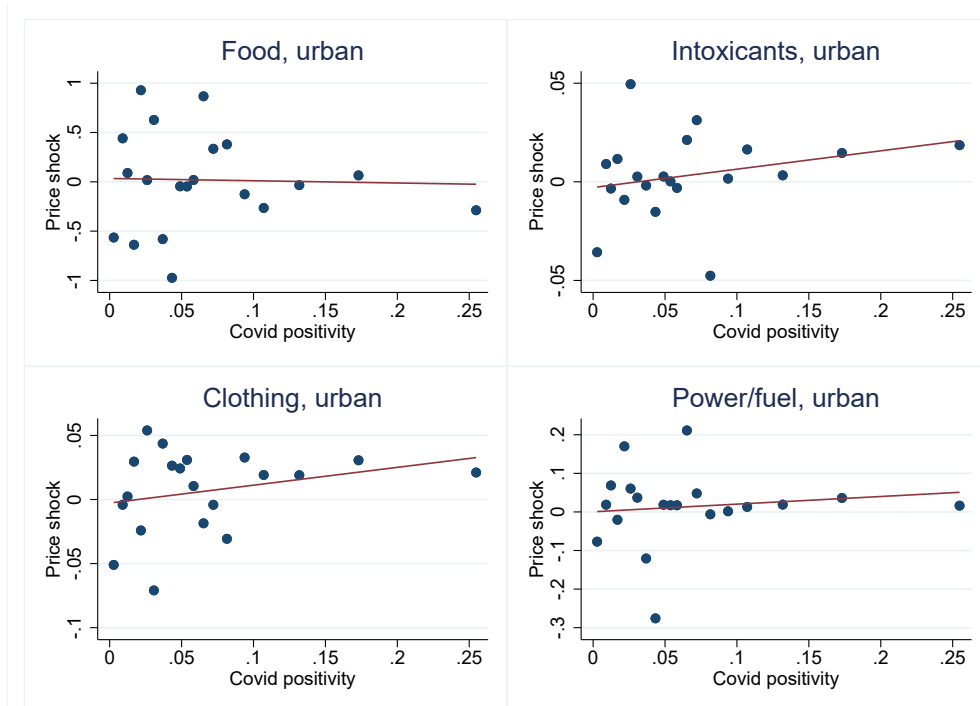


Figure 15: Relationship between Price Shocks and COVID positivity rates

Panel A: Rural and urban areas



Panel B: Only urban areas



Note. Figure reports binscatter of price shocks, across urban and rural areas separately, against estimates of the COVID-19 positivity rate, for different categories. Covid data are from www.covid19India.org.

the price and nominal income shocks were negative, but intoxicants are a somewhat inferior good ($\beta = -0.02$.) Fuel consumption lay somewhere between food and intoxicants: share rose, but mainly because income fell and fuels is mildly inferior ($\beta = -0.07$). Price somewhat offset nominal income's effect, but it was a minor offset.

When the economy began to recover in August 2020 (Figure 14B), food consumption continued to be elevated. Since we have separate overall and food CPI for this month, we can see that, although price overall played a small role, it was purchasing power that did most of the work. An increase in overall inflation reinforced the decline in nominal income, raising consumption of food. What is perhaps most notable about this is that relative prices played such a small role. Somehow, through one of the most severe lockdowns around the world, the country was able to maintain the relative price of food.

Table 7: Covid prevalence and price shocks

	(1)	(2)	(3)	(4)
	Positivity rate (urban)	Positivity rate (urban & rural)	Cases per capita (urban)	Cases per capita (urban & rural)
Food	-0.231 (1.789)	-0.285 (1.100)	-20.958 (59.940)	-15.639 (37.013)
Intoxicants	0.093 (0.077)	0.093 (0.064)	1.210 (2.593)	1.280 (2.149)
Clothes/shoes	0.139 (0.119)	0.128 (0.091)	2.035 (3.987)	1.783 (3.065)
Power/fuel	0.198 (0.409)	0.209 (0.271)	1.651 (13.702)	2.454 (9.128)
$\beta_{Food} = \beta_{Intoxicants}$	0.855	0.728	0.709	0.644
$\beta_{Food} = \beta_{Clothes/shoes}$	0.841	0.713	0.710	0.645
$\beta_{Food} = \beta_{Power/fuel}$	0.771	0.586	0.647	0.553
N	195	383	195	383
R^2	0.000	0.000	0.001	0.000

Notes: Observations are at the state \times community type \times month level. Significance levels: * 10% ** 5% *** 1%.

Price shocks and COVID cases. A natural question is whether areas with greater price shocks were areas harder hit by COVID cases. We back out price shocks by location (state \times community type) and product for April 2020. We measure COVID in two ways: number of confirmed cases divided by number of tests (positivity rate) or simply the number of confirmed cases in that month. The positivity rate has the advantage that it accounts for the low level of testing, especially early in the pandemic.

We present suggestive graphical evidence in Figure 15 that, outside of food, areas with greater COVID positivity rates appear to have greater positive price shocks in April 2020 (relative to April 2019). This is only suggestive, however. Regressions of price shocks on COVID cases do not reveal significant differences in price shocks for food versus other categories. Table 7 reports coefficients from 16 regressions. In each we regressed implied price shocks for a given product category on measures of COVID by location: $\Delta p_{j,k,t} = \lambda_{k,t} d_{j,k,t} + e_{j,k,t}$. Estimates of $\lambda_{k,t}$ are reported, but none are significantly different from one another.

Nonetheless, if we take the estimated correlation seriously, there are two possible ex-

planations for the findings. One is that the lockdown was more severe—at least with respect to non-food items—in areas that were harder hit by COVID. Another is that there is a shadow price associated with COVID for non-essential items, i.e., individuals' perceived risk of getting infected.

5 Conclusion

We draw four conclusions from our analysis. First, the economic shock of COVID in India was harsh, but short-lived. That said, there appear to be meaningful disparities in how badly household incomes fell and in how much households had recovered as of October 2020. Those trends should be followed to see if COVID increased income inequality in India.

Second, workers tried to smooth their income by switching occupations. However, this switching may have had spillover effects on other occupations and changed the distribution of income losses across occupations. Further work is required to develop counterfactuals—what workers would have earned had they not switched occupations during the height of the COVID shock—to better estimate the degree of loss reallocation.

Third, the COVID-induced economic shock generated less disparity in consumption than income. Our analysis suggests that households were able to insure idiosyncratic shocks as well as before COVID. An intriguing possibility is that the precautionary savings of higher income households made it easier for lower income households to borrow to buffer consumption. It may be an interesting feature of aggregate supply shocks, or at least pandemic induced supply shocks, that they facilitate inter-household insurance through, e.g., credit channels. This possibility can be explored using the same Engel curve decomposition used to understand the relative effect of income and price changes on the share of household budgets spent in different product categories.

Fourth, the change in budget shares across product categories suggests that the relative price of food did not increase during the pandemic, indicating that the lockdown may have been somewhat successful at discriminating between essential and non-essential services. Other product categories, such as clothing, did experience increases in relative prices. Future work will explore whether geographic variation in the relative price shock is correlated with infection risk and non-pharmaceutical interventions.

References

- Abraham, R. and A. Shrivastava (2019). How comparable are india's labour market surveys? Report, Centre for Sustainable Employment, Azim Premji University.
- Aguiar, M. and M. Bils (2015). Has consumption inequality mirrored income inequality? *American Economic Review* 105(9), 2725–56.
- Almås, I. (2012, April). International income inequality: Measuring ppp bias by estimating engel curves for food. *American Economic Review* 102(2), 1093–1117.
- Alstadsæter, A., B. Bratsberg, G. Eielsen, W. Kopczuk, S. Markussen, O. Raaum, and K. Røed (2020). The first weeks of the coronavirus crisis: Who got hit, when and why? evidence from norway. Report 0898-2937, National Bureau of Economic Research.
- Atkin, D., B. Faber, T. Fally, and M. Gonzalez-Navarro (2020). Measuring welfare and inequality with incomplete price information. *National Bureau of Economic Research Working Paper Series No. 26890*.
- Bardhan, P. K. (1983). Labor-tying in a poor agrarian economy: A theoretical and empirical analysis. *The Quarterly Journal of Economics* 98(3), 501–514.
- Bertrand, M., R. Dizon-Ross, K. Krishnan, and H. Schofield (2020). Employment, income and consumption in india during and after the lockdown: A v-shape recovery? Report, Rustandy Center for Social Sector Innovation Blog.
- Bertrand, M., K. Krishnan, and H. Schofield (2020). How are indian households coping under the covid-19 lockdown? 8 key findings. *Rustandy Center for Social Sector Innovation Blog* 11.
- Beyer, R., T. Jain, and S. Sinha (2020). Lights out? covid-19 containment policies and economic activity. Report.
- Bhoi, B., H. Shekhar, and I. Padhi (2020). Rural-urban inflation dynamics in india. *RBI Bulletin December*.
- Binswanger, H. and M. Rosenzweig (1993). Wealth, weather risk and the composition and profitability of agricultural investments. *Economic Journal* 103(416), 56–78.
- Carrillo-Tudela, C. and L. Visschers (2020). Unemployment and endogenous reallocation over the business cycle. Report.
- Chaudhuri, S. and C. Paxson (2002). Smoothing consumption under income seasonality: buffer stocks vs. credit markets.
- Chen, H., W. Qian, and Q. Wen (2020). The impact of the covid-19 pandemic on consumption: Learning from high frequency transaction data. *Available at SSRN* 3568574.
- Chetty, R. and A. Looney (2007). *Income risk and the benefits of social insurance: Evidence from Indonesia and the United States*, pp. 99–121. University of Chicago Press.


- Chodorow-Reich, G. and J. Wieland (2020). Secular labor reallocation and business cycles. *Journal of Political Economy* 128(6), 2245–2287.
- Cochrane, J. H. (1991). A simple test of consumption insurance. *Journal of political economy* 99(5), 957–976.
- Costa, D. L. (2001). Estimating real income in the united states from 1888 to 1994: Correcting cpi bias using engel curves. *Journal of political economy* 109(6), 1288–1310.
- Deaton, A. (2021). Covid-19 and global income inequality. *National Bureau of Economic Research Working Paper Series No. 28392*.
- Deaton, A. and J. Muellbauer (1980). An almost ideal demand system. *The American Economic Review* 70(3), 312–326.
- Deb, P., D. Furceri, J. D. Ostry, and N. Tawk (2020). The economic effects of covid-19 containment measures.
- Dercon, S. (2002). Income risk, coping strategies, and safety nets. *The World Bank Research Observer* 17(2), 141–166.
- Deshpande, A. (2020). The covid-19 pandemic and lockdown: First effects on gender gaps in employment and domestic work in india? Report, Working Paper 30, Ashoka University, Department of Economics, revised June.
- Dhingra, S. and S. J. Machin (2020). The crisis and job guarantees in urban india. Report, IZA DP No. 13760.
- Egger, D., E. Miguel, S. S. Warren, A. Shenoy, E. Collins, D. Karlan, D. Parkerson, A. M. Mobarak, G. Fink, C. Udry, M. Walker, J. Haushofer, M. Larrebourg, S. Athey, P. Lopez-Pena, S. Benhachmi, M. Humphreys, L. Lowe, N. F. Meriggi, A. Wabwire, C. A. Davis, U. J. Pape, T. Graff, M. Voors, C. Nekesa, and C. Vernot (2021). Falling living standards during the covid-19 crisis: Quantitative evidence from nine developing countries. *Science Advances* 7(6), eabe0997.
- Frankenberg, E., J. P. Smith, and D. Thomas (2003). Economic shocks, wealth, and welfare. *Journal of Human Resources* 38(2), 280–321.
- French, E. and C. Taber (2011). *Identification of models of the labor market*, Volume 4, pp. 537–617. Elsevier.
- Google LLC (2021). Google covid-19 community mobility reports.
- Groes, F., P. Kircher, and I. Manovskii (2015). The u-shapes of occupational mobility. *The Review of Economic Studies* 82(2), 659–692.
- Guerrieri, V., G. Lorenzoni, L. Straub, and I. Werning (2020). Macroeconomic implications of covid-19: Can negative supply shocks cause demand shortages? Report 0898-2937, National Bureau of Economic Research.

- Hale, T., S. Webster, A. Petherick, T. Phillips, and B. Kira (2020). Variation in government responses to covid-19. Report, Blavatnik school of government working paper.
- Hamilton, B. W. (2001). Using engel’s law to estimate cpi bias. *American Economic Review* 91(3), 619–630.
- Hoseini, M. and T. Beck (2020). Emergency loans and consumption—evidence from covid-19 in iran. *Covid Economics* 45(28), 111–146.
- Kambourov, G. and I. Manovskii (2008). Rising occupational and industry mobility in the united states: 1968–97. *International Economic Review* 49(1), 41–79.
- Kambourov, G. and I. Manovskii (2009a). Occupational mobility and wage inequality. *The Review of Economic Studies* 76(2), 731–759.
- Kambourov, G. and I. Manovskii (2009b). Occupational specificity of human capital. *International Economic Review* 50(1), 63–115.
- Kochar, A. (1999). Smoothing consumption by smoothing income: hours-of-work responses to idiosyncratic agricultural shocks in rural india. *Review of Economics and Statistics* 81(1), 50–61.
- Krueger, D. and F. Perri (2006). Does income inequality lead to consumption inequality? evidence and theory1. *The Review of Economic Studies* 73(1), 163–193.
- Lee, K., H. Sahai, P. Baylis, and M. Greenstone (2020). Job loss and behavioral change: The unprecedented effects of the india lockdown in delhi. *University of Chicago, Becker Friedman Institute for Economics Working Paper* (2020-65).
- Logan, T. D. (2009). Are engel curve estimates of cpi bias biased? *Historical Methods: A Journal of Quantitative and Interdisciplinary History* 42(3), 97–110.
- Malani, A., M. Mohanan, C. Kumar, J. Kramer, and V. Tandel (2020). Prevalence of sars-cov-2 among workers returning to bihar gives snapshot of covid across india. *medRxiv*, 2020.06.26.20138545.
- Malani, A., S. Ramachandran, V. Tandel, R. Parasa, S. Sudharshini, V. Prakash, Y. Yogananth, S. Raju, and T. S. Selvavinayagam (2021). Sars-cov-2 seroprevalence in tamil nadu in october-november 2020. *medRxiv*, 2021.02.03.21250949.
- Mohanan, M., A. Malani, K. Krishnan, and A. Acharya (2020). Prevalence of covid-19 in rural versus urban areas in a low-income country: Findings from a state-wide study in karnataka, india. *medRxiv*, 2020.11.02.20224782.
- Morduch, J. (1995, September). Income smoothing and consumption smoothing. *Journal of Economic Perspectives* 9(3), 103–114.
- Moscarini, G. and K. Thomsson (2007). Occupational and job mobility in the us. *The Scandinavian Journal of Economics* 109(4), 807–836.

- Moser, C. O. N. (1998). The asset vulnerability framework: Reassessing urban poverty reduction strategies. *World Development* 26(1), 1–19.
- Nakamura, E., J. Steinsson, and M. Liu (2016, July). Are chinese growth and inflation too smooth? evidence from engel curves. *American Economic Journal: Macroeconomics* 8(3), 113–44.
- Nakamura, L. (1995). Is u.s. economic performance really that bad? Working Papers 95-21, Federal Reserve Bank of Philadelphia.
- Paulson, A. (1994). Insurance motives for migration: Evidence from thailand.
- Pinto, A., A. Bhowmick, and R. Adlakha (2020). How did india’s rural economy fare through the covid-19 lockdown and the re-opening? Report, World Bank.
- Reserve Bank of India (2020). Annual report 2019–2020. Report, Reserve Bank of India.
- Rose, E. (2001). Ex ante and ex post labor supply response to risk in a low-income area. *Journal of Development Economics* 64(2), 371–388.
- Rosenzweig, M. R. and O. Stark (1989). Consumption smoothing, migration, and marriage: Evidence from rural india. *Journal of political Economy* 97(4), 905–926.
- Townsend, R. M. (1994). Risk and insurance in village india. *Econometrica* 62(3), 539–591.
- Vyas, M. (2020a). Cphs execution during the lockdown of 2020. Report, Centre for Monitoring Indian Economy.
- Vyas, M. (2020b). The questionnaire. Report, Centre for Monitoring Indian Economy.
- Vyas, M. (2020c). Sample survival response rate. Report, Centre for Monitoring Indian Economy.
- Vyas, M. (2020d). Survey design and sample. Report, Centre for Monitoring Indian Economy.
- Worldometer (2021). Covid-19 coronavirus pandemic. Report.
- Young, A. (2012). The african growth miracle. *Journal of Political Economy* 120(4), 696–739.


Appendix

Figure A1: Implied urban and rural wages conditional on positive hours, by occupational category.




Note. Includes males and females aged 15 or above. Wages are coefficients from a regression of per capita household income (includes wages plus household level business income) on hours, run separately by occupation and location. We report median implied wages across locations, by occupation. Dashed vertical lines in January 2020, March 2020 and June 2020 indicate the month of first case (blue), the month the national lockdown started (red) and the month the national lockdown ended (green).

Figure A2: Implied urban and rural wages conditional on positive hours, by occupational category.



Note. Includes all males (not females) aged 15 or above. Wages are coefficients from a regression of per capita household income (includes wages and household level business income) on hours, run separately by occupation and location. We report median implied wages across locations, by occupation. Dashed vertical lines in January 2020, March 2020 and June 2020 indicate the month of first case (blue), the month the national lockdown started (red) and the month the national lockdown ended (green).

Figure A3: Implied urban and rural wages conditional on positive hours, by occupational category.



Note. Includes all males (not females) aged 15 or above. Wages are coefficients from a regression of wage income (not household income) on hours, run separately by occupation and location. We report median implied wages across locations, by occupation. Dashed vertical lines in January 2020, March 2020 and June 2020 indicate the month of first case (blue), the month the national lockdown started (red) and the month the national lockdown ended (green).