# Credit Where Credit Is Withdrawn: Weather Shocks and the Household Coping Portfolio in India

Anand Mathew* Sankalp Mathur†

November 2026

**Abstract**

In September 2018 the default of Infrastructure Leasing and Financial Services, an infrastructure financier rated AAA weeks earlier, cut the wholesale funding of India's non-bank lenders and bound hardest in the districts that relied on them. We use this dislocation to ask what a contraction in formal credit supply does to the way the households of a district absorb a dry shock. In a panel of 210,560 households observed from 2014 to 2019, set beside 483 million loan records aggregated to the same districts, a dry shock takes the same share of income from exposed occupations whether the district's exposure to the contraction was higher or lower, and leaves total spending close to unchanged in both cases, while the composition of the basket differs. Where exposure was lower, the cuts begin with recreation and restaurant meals and stop before health, education and a daily staple such as milk. Where it was higher, the same shock reaches all three. On the borrowing margins, fewer households hold a non-bank loan, more hold one from a moneylender, though imprecisely, and among households that lose formal access about as many stop borrowing altogether as keep or find an informal source. Few households borrow from a non-bank directly, so the effect runs through the district's credit market and not only through the household's own lender. Credit does not stop a dry shock from being felt, and it determines how far down the coping portfolio a household has to go.

Keywords: consumption smoothing, credit supply, weather shocks, shadow banks, India
JEL: O12, O16, G21, Q54, D14

*University of Florida, corresponding author. All errors are our own.
†CAFRAL, Reserve Bank of India. The views expressed here are those of the authors and do not necessarily represent those of CAFRAL or the Reserve Bank of India.

# 1 Introduction

What does formal credit do for a household whose income has just fallen? The Indian retail credit market of the past decade poses the question at a scale that the village studies which first asked it could not. Between 2015 and 2019, formal lenders originated 483 million retail loans worth 105 trillion rupees across 640 districts, and non-bank lenders raised their share of the rupees lent from 22 percent to 29.1 percent by the middle of 2018. The shocks we study are the months that a district's own history marks as unusually dry or unusually wet, and since every district meets them, the question is not whether households face such months but how they absorb them. A sizeable literature shows that giving households access to finance helps them absorb such shocks, whether the access arrives through a guaranteed loan, a savings group, a digital lender or a mobile wallet (Lane, 2024; Demont, 2020; Suri et al., 2021; Jack and Suri, 2014). Much less is known about the reverse event, in which a lender that households had come to rely on withdraws, and on the credit side of the question the household has so far been observed only in the level of its consumption, never in the portfolio of responses it chooses from. A credit file records what lenders did and nothing of what the household did instead.

When the rains fall short in an Indian district, the households living there face a gap between what they have committed to spend and what they are about to earn, and the gap has to be closed somehow. A household can borrow from a bank or a non-bank lender, borrow from a moneylender or a relative, sell an animal or a piece of land, stop putting money into gold, put another member into the labour market, or buy less. These options are not interchangeable, since some reverse when the rain returns and some may not, and a household that sells a productive asset in a bad year may never recover the position (Carter and Barrett, 2006). We call the set of these margins the coping portfolio, and the question of this paper is how far along it a household must travel, and whether formal credit changes the distance.

Where the response shows is itself part of the answer. Weather shocks cut income in India and household consumption barely moves, a null found often enough to look like a fact about insurance (Townsend, 1994; Morduch, 1995). We argue that it is in part a fact about the outcome variable. Households change what they buy before they change how much, so a regression on the level of spending looks for a change in scale that never occurs, while the composition of the basket records the adjustment. The pattern has been documented in rural India and read as successful smoothing (Malaeb et al., 2026). What has not been done, and what this paper does, is tie it to the supply of credit. Cramer et al. (2025), working with the same bureau records we use, find that non-bank lenders increase originations after an extreme weather month, and by more than banks do, but a credit file cannot say what happens to the household on the other side of the transaction. The two sides have to be observed together, for the same households at the same time, under a change in credit supply that households did not choose. The experiments that established the value of access add credit at the scale of a village or a kiosk. A contraction at the scale of a country cannot be randomised, and India's non-bank sector provided one in

2018.

In September 2018, Infrastructure Leasing and Financial Services (IL&FS), a conglomerate that financed roads and power projects and had carried the highest credit rating until weeks before, defaulted on its obligations. The defaults drove debt mutual funds, the main buyers of non-bank commercial paper, out of the market, so that rollover rates on that paper fell from above 95 percent to below 10 percent within three months (Sengupta et al., 2022), and non-bank lenders that had funded long-duration retail loans with short-duration paper could no longer roll it. The non-bank and housing-finance share of the rupees lent to Indian households had risen from 22 percent at the start of 2015 to 29.1 percent in May 2018, and by the end of 2019 it stood at 20.6 percent (Figure 1).

The paper combines three sources, the first of which is the CMIE Consumer Pyramids Household Survey, from which we observe, for 210,560 households in 525 districts from January 2014 to December 2019, spending on over a hundred and fifty items, income by source, borrowing by source and by purpose, asset sales, gold saving, and who in the household works. From TransUnion CIBIL, India's largest credit bureau, we observe the formal retail lending that reaches the bureau, 483 million originations across 640 districts from 2015 to 2019, by month, product and lender type. And from the Standardised Precipitation and Evapotranspiration Index we measure, for every district and month, whether the water balance fell into the dry or the wet tail of that district's own history.

We identify the household consequences of the contraction from the fact that it bound differently across places. A district's reliance on non-bank lenders is measured as the share of loans originated by non-bank and housing-finance lenders over the twelve months before the first sign of trouble, June 2017 to May 2018, and is held fixed thereafter, as Breza and Kinnan (2021) hold district exposure fixed on the eve of the microfinance crisis, so it cannot respond to the default or to any weather that follows. In the top tercile of this share, non-bank lending fell by 8 log points more than in the bottom tercile between the first half of 2018 and the end of 2019 (Panel B of Figure 1), and by 12 percent in the regression of Table 1. Before the default the weather response of households in the two kinds of district had not been diverging (Figure 2), though their lending had, a caveat the paper carries throughout. We then ask what the same weather shock does to households in the first kind of district compared with the second, before and after the default. The comparison is a triple difference that absorbs any standing tendency of non-bank-reliant districts to respond differently to weather and any post-default divergence in those districts unrelated to weather, and it is estimated within states, so what remains is the change at the default in how a dry shock reaches each margin of the portfolio where exposure to the contraction was higher, relative to where it was lower, and the estimand is a district-level one, measuring how the households of a district absorb a dry shock once formal credit in the district has contracted, through their own lender and through every local channel the contraction touches. The comparison is disciplined by two features of the setting. The default originated in

infrastructure finance, so nothing about it responds to rainfall in any district, and salaried and retired households, whose incomes do not depend on the monsoon, sit in the same districts and months and serve as placebos throughout.

The weather shock on its own produces three regularities. First, weather reaches households through occupation, and the two tails reach different occupations. A deficit works through the crop, so agricultural labourers lose 6.2 percent of household income and small farmers 4.8 percent. A surplus works through stopped outdoor work, so construction labourers lose 5.5 percent and the outdoor self-employed 4.8 percent, while salaried households in the same districts and months lose nothing under either tail. Second, the level of spending does not respond while its composition does. Total expenditure moves by less than one percent in either tail and is indistinguishable from zero, while the discretionary share of the budget falls by 0.68 percentage points on a base of 26.7 percent after a dry shock. Third, within the basket the cuts run from the most postponable spending to the least. Recreation and restaurant meals fall by double digits, while cereals and power do not fall.

The credit contraction sets how far down that ordering a household must go, and the comparison begins with what does not differ. The same dry shock takes the same share of income from the exposed occupations in both kinds of district, 0.2 percent apart with a standard error eight times that gap, and total spending moves by the same small amount in both, so what follows is a difference in how a loss is absorbed and not in its size. Where exposure was lower, a dry shock moves the discretionary share by 0.42 percentage points and leaves health, education and a staple such as milk untouched. Where it was higher, the same shock moves the share by 0.94 points and cuts health spending by 3.5 percent more, education by 4.4 percent more and milk by 2.7 percent more, an index of the three by 3.2 percent. The discretionary rows roughly double, which a larger shock would also have produced. The protected rows move from zero, which a shock of the same size would not have produced, and income fell by the same amount in both regimes within the precision of the data. This distinction between a deeper cut and a wider one is the paper's central result.

On the borrowing margins the household side moves as the supply side says it should. Where exposure was higher, 0.66 percentage points fewer households hold a non-bank loan after a dry shock than where it was lower, a fifth of the margin's base, bank borrowing is imprecise and neither confirms nor contradicts the supply-side evidence that banks held their lending, the share borrowing from a moneylender rises by 0.64 points more, with a standard error of two thirds that size, and the instalments households pay on their loans, the one rupee measure of debt the survey carries, fall by 6.7 percent more. Classifying each household-wave by the joint state of its formal and informal borrowing, we find that a dry shock where exposure was higher raises the share of formal borrowers who hold no loan from any source a wave later by 2.7 percentage points and the share who keep or find an informal one by 2.2 points, with a standard error of 1.8 on the difference, so losing a formal lender is about as often rationing as redirection. Because

the survey records who borrows but never how much, these are statements about households and not about rupees. The direct margin cannot carry the spending result on its own, since only 3.1 percent of household-waves hold a non-bank loan, so most of what the portfolio tables record must run through the district and not through the household's own lender. Breza and Kinnan (2021) found the same when microcredit contracted, with consumption falling in households that had never borrowed, and we read our estimates in the same way, as the effect of a contraction in a district's formal credit on how its households absorb a dry shock, with the household's own lender as one channel among several.

A design that leans this heavily on placebo comparisons should be tested where it ought to fail, and most of the tests return nothing. The salaried placebo is null on income and on the four main margins, future weather does not predict current spending, district originations of large loans do not respond to dry months, and weather shocks do not push households out of the panel. The headline estimate is identified within states and is significant at ten percent, the pooled design that uses the cross-state variation as well gives an estimate a quarter larger that is significant at five percent, and the estimate keeps its size when the states that waived farm loans in 2018 are excluded and under spatial and state-clustered inference, where its $p$-value lies between ten and fifteen percent, while it is significant at ten percent under the permutation of exposure across all districts. Some results do not fit the account, since loan counts rise where rupees fall, overdue shop credit moves the wrong way, one post-default block dips, and the wet tail shows no credit interaction. Two further results bear on the design itself, a wet-tail response among retired households and a pre-period in the event study too imprecise to rule out drift of the size of the effect, and each of these is discussed beside the result it concerns.

The paper speaks first to the literature on how households in poor agrarian economies insure income risk. Townsend (1994) found in the ICRISAT villages substantial but incomplete risk sharing, which has since been the standard reading (Morduch, 1995), and the mechanisms have been traced one at a time, savings out of transitory income (Paxson, 1992), bullock sales at real cost to productive capital (Rosenzweig and Wolpin, 1993), informal loans whose repayment depends on the shocks of both parties (Udry, 1994), and livestock that buffers much less than theory predicts (Fafchamps et al., 1998; Kazianga and Udry, 2006). For India, Jayachandran (2006) established that productivity shocks move rural wages sharply, and Kaur (2019) showed that nominal rigidity pushes the adjustment into employment, and the practice of defining a shock in the tails of a district's own rainfall history comes from that work (Shah and Steinberg, 2017; Corno et al., 2020), which we carry to the monthly frequency of the lending data. Closest in design are the studies that pair exogenous access to finance with shocks, a guaranteed emergency loan after a flood (Lane, 2024), randomised self-help group access against rainfall shocks (Demont, 2020), a credit-score cutoff for digital loans (Suri et al., 2021) and the spread of mobile money (Jack and Suri, 2014), all of which find that access blunts the effect of a shock on consumption or on the expenses households forgo. We study the removal of credit where they

study its addition, at the scale of a country rather than a village, with the supply shock observed directly in the bureau, and with the whole portfolio observed at once, so that a household that stops borrowing and starts selling is seen doing both.

On the credit-supply side the gap in the literature is a specific one. Cramer et al. (2025) interact the same weather shock with the IL&FS date and find that non-bank lending stops responding to weather after the default, but they observe no household. Breza and Kinnan (2021) study a contraction in microcredit across Indian districts and find lower wages, earnings and consumption, but rainfall enters their analysis only as a control. This paper supplies the household counterpart of the first and the shock counterpart of the second. It also draws on the literature that identifies credit effects from shocks to supply (Khwaja and Mian, 2008; Chodorow-Reich, 2014; Huber, 2018), on the studies of the IL&FS episode as an event in financial stability (Acharya et al., 2013; Kulkarni et al., 2023; Bhardwaj and Javadekar, 2024), on the work that attributes the growth of non-bank lenders in the United States mostly to regulatory arbitrage (Buchak et al., 2018; Irani et al., 2021; Gopal and Schnabl, 2022), and on the work on India's other buffers, the employment guarantee (Imbert and Papp, 2015), transfer delivery (Muralidharan et al., 2016) and the farmer income transfer (Ghosh and Vats, 2024), whose positive, permanent and predictable transfer finances investment where our negative, transitory and unpredictable shock moves the composition of consumption.

Section 2 describes the default, its transmission to district credit supply, the considerations that shape identification, and the predictions the coping portfolio makes. Section 3 describes the data and the empirical strategy and presents the event study. Section 4 reports the effects of the default on lending and of the weather on lending and on earnings, the results on every margin of the portfolio, and the robustness checks. Section 5 interprets the estimates and states their limits, and Section 6 concludes.

## 2 Setting and Empirical Predictions

On 4 September 2018, Infrastructure Leasing and Financial Services defaulted on a short-term loan from the Small Industries Development Bank of India, and over the following weeks the market for the short-term paper that funded India's non-bank lenders closed (Sengupta et al., 2022). The group was an infrastructure financier, not a retail lender, and what matters for the paper is that the lenders households did borrow from lost their funding, and that they lost it for reasons that had nothing to do with the weather in any district.

### 2.1 The IL&FS Default

IL&FS was founded in 1987 and registered with the Reserve Bank of India as a core investment company. It financed and built infrastructure through 169 legal entities, many of them special purpose vehicles that owned a single asset, and a non-bank finance subsidiary that lent to those

vehicles (Sengupta et al., 2022; Bhardwaj and Javadekar, 2024). By September 2018 the group's debt stood near 910 billion rupees, with a ratio of debt to equity of about seventeen, three times the sector average, and its instruments carried the highest rating from most agencies and were widely held by mutual funds, pension funds and corporate treasuries (Sengupta et al., 2022; Bhardwaj and Javadekar, 2024). The first sign of trouble came in June 2018, when a transport subsidiary failed to repay inter-corporate deposits, the group's financial arm defaulted on its commercial paper in August, and on 4 September the parent defaulted on a loan from the Small Industries Development Bank, after which the agencies that had rated the group AAA in July cut it to junk within weeks and the government took control of its management (Sengupta et al., 2022; Bhardwaj and Javadekar, 2024; Cramer et al., 2025).

Three features of the failure matter for what follows. Its causes were internal to the group, delays and cost overruns on infrastructure projects, a maturity mismatch and failures of governance, and they had nothing to do with the state of the rural economy or with the rainfall in any district (Bhardwaj and Javadekar, 2024; Cramer et al., 2025). It was not anticipated, since the group held the highest credit rating until weeks before it failed, and the equity of the largest retail non-bank lender, which had minimal exposure to infrastructure, fell by about a quarter in the two months after the default (Cramer et al., 2025). And it was large enough to reprice the funding of the whole sector and not of one firm, which is the subject of the next subsection.

### 2.2 Transmission to District Credit Supply

Non-bank lenders in India fund long-duration retail assets, vehicle loans, loans against property, consumer durables and housing, with short-duration wholesale liabilities, and debt mutual funds are the marginal buyers of that paper, holding about four fifths of the commercial paper issued by non-banks in March 2019 (Government of India, 2020). Acharya et al. (2013) had described the fragility of this structure years before the default. When the IL&FS paper defaulted, investors in liquid debt funds ran, the funds stopped rolling over the commercial paper of other non-banks, rollover rates fell from above 95 percent to below 10 percent, and within three months the supply of funds to thousands of non-bank lenders had dried up (Government of India, 2020; Sengupta et al., 2022). The lenders with the largest short-term mismatches lost the most funding and could not replace it with bank credit or debentures in the short run (Kulkarni et al., 2023), while banks with exposure to IL&FS cut their lending to unrelated non-banks by nearly half (Bhardwaj and Javadekar, 2024).

Panel A of Figure 1 shows the consequence in the origination data. The non-bank and housing-finance share of the rupees lent rises from 22 percent in January 2015 to 29.1 percent in May 2018 and stands at 28.4 percent in July, turns at the default, falls to 24.0 percent by October 2018, and ends the panel at 20.6 percent in December 2019. The share of loan counts holds up through most of 2019 and falls only at its end, from 36.1 percent in July 2018 to 27.9 percent in December 2019, so in counts the contraction arrived late, and the first year's fall in the rupee

share came from the rupees per loan of the whole market, as banks and financial-technology lenders added small loans. The Reserve Bank lowered the risk weights on bank lending to non-banks and stressed lenders sold loan portfolios to banks, which kept some credit flowing but did not restore originations, and the contraction was felt most acutely in the first half of the financial year beginning April 2019 (Sengupta et al., 2022). The contraction continues past December 2019, but the household panel does not, so everything estimated here stops where the survey stops.

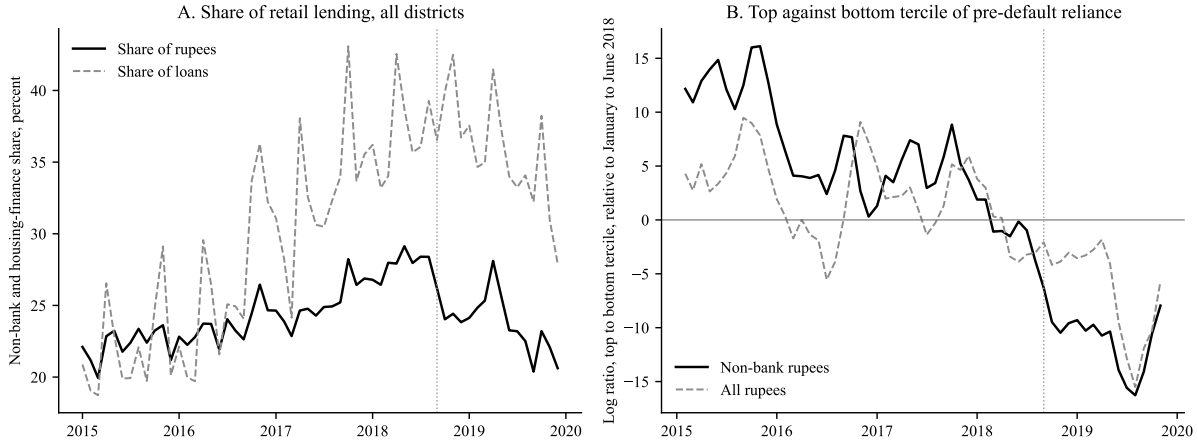



**Figure 1:** The Non-Bank Contraction after September 2018

Panel A plots the non-bank and housing-finance share of the rupees and of the loans originated each month across 640 districts, 2015 to 2019. Panel B plots the log of the ratio of rupees originated in the top tercile of pre-default non-bank reliance to rupees originated in the bottom tercile, for non-bank and housing-finance lenders and for all lenders, three-month centred moving averages expressed relative to the mean of January to June 2018. The dotted line marks September 2018. Non-bank gold loans and the 2017 burst of farm credit cards reported by other lenders are left out throughout, as Appendix C explains.

The contraction was national, so by itself it is absorbed by any time effect. What identifies its household consequences is that it bound differently across places, and Panel B of Figure 1 shows the variation the rest of the paper uses, as the ratio of lending in the top tercile of pre-default non-bank reliance to lending in the bottom tercile. The ratio of non-bank rupees falls by 8 log points between the first half of 2018 and the end of 2019, with the break at the default plainly visible, and the ratio of all rupees falls by 6 log points, more than the non-bank share of the market would produce on its own, so other lenders did not fill the gap in rupees and fell a little themselves. The panel also shows the caveat that Section 2.3 takes up. The ratio had been drifting down since 2015, because non-bank lenders were expanding faster in low-reliance districts from a low base, so the two groups were not on parallel paths in lending before the default, and the break at September 2018 sits on top of that drift.

### 2.3 Identifying the Role of the Contraction

We use exposure to the funding shock as a change in the supply of formal credit that households did not choose, and the account above suggests six considerations for the design.

First, districts that relied on non-bank lenders before the default differ from others, and Appendix Table A3 shows how on the districts that enter the estimates. The top tercile is more urban, with 34 percent of its population in towns against 26 percent in the bottom tercile, a third more populous, and its households borrowed about three quarters more non-bank rupees per person in the year before the default, while the bottom tercile, served by banks, had nearly twice the loans per person, most of them small farm loans, and about the same rupees per person overall. The bottom tercile also had one percentage point more dry-tail months over the sample, and after the default the difference widens, since the dry months of 2019 fell on the south and the north-east, where reliance on non-banks was low, so that a dry month was recorded in 21 percent of post-default months in the bottom tercile against 14 percent in the top. The weather is exogenous and the imbalance costs precision rather than consistency, but the identifying cell is thinner than the sample suggests, and Appendix Table A2 counts the district-blocks in every cell of the design. Households in exposed districts may respond to a dry shock differently for all of these reasons. The exposure measure is therefore frozen before the default and the specification carries every lower-order interaction, so that the identifying variation is the change at the default in the weather response of exposed districts relative to unexposed ones, and the identifying assumption is that no such change would have occurred for reasons unrelated to the default. The main specification lets the weather response and its change at the default differ by state, so that the estimate comes from the 10 states that hold districts in both terciles, and Section 4.3 reports the pooled design that uses the cross-state variation as well, along with every district characteristic in the balance table interacted with the shock and the post-default indicator.

Second, the two groups of districts were not on parallel paths in lending before the default, as Panel B of Figure 1 shows. Section 4.1 therefore reports the lending effect with and without district-specific trends, and the household estimates rest on the weather response and not on the level of lending, which the trend in lending does not touch directly.

Third, the exposure measure is a baseline share, and Goldsmith-Pinkham et al. (2020) show that a design built on such a share rests on the share being as good as randomly assigned conditional on the controls, the same assumption whether or not the share is called an instrument, and we do not call it one. We report its construction, distribution and correlates, treat the tercile contrast as a difference in differences with a continuous dose behind it in the sense of Callaway et al. (2024), and report the dose itself, a share measured two years earlier, a second exposure built from vehicle and mortgage lending, the products whose lenders depended most on commercial paper, and a share that includes gold loans, in Section 4.3. Instrumenting reliance with early adoption of digital-payments infrastructure, following Dubey and Purnanandam (2023), would

relax the assumption, and it needs branch-level deposit data that we do not hold.

Fourth, the post-default window is short, since the household panel ends in December 2019, sixteen months after the default, which gives four four-month blocks after the event, and whether the gap that opens then stabilises or widens is beyond these data.

Fifth, other things happened in 2018 and 2019, and the balance table says where they landed. Karnataka waived farm loans in July 2018, Madhya Pradesh, Chhattisgarh and Rajasthan did so in the week after their December 2018 elections, a cash transfer to landowning farmers began paying in February 2019 (Ghosh and Vats, 2024), and the national hospital insurance scheme for poor households was launched on 23 September 2018, three weeks after the default. Each relaxes the constraint of some households after a dry shock, in the post-default period, and farming households sit disproportionately in the bottom tercile, so each could produce a difference between the terciles with no funding shock behind it. None of them is a weather programme, so the state-by-block effects absorb their levels, and what would bias $\theta$ is a change at the default in the weather response of bank-served farm districts that they produce, which the within-state design and the exclusions of Section 4.3 address. The insurance scheme bears on the health row in particular, and the cut it would have to explain appears in education and in a staple, milk, as well, which no health scheme touches. The slowdown in vehicle and property markets through 2019 cuts the other way, since it lowered the demand for the loans non-banks make in the urban districts where they made them. The state-by-block effects absorb the national timing of all four, the main specification lets the weather response and its change at the default differ by state, Section 4.3 reports the estimates with the waiver states excluded and with farming households excluded, and Appendix E splits the estimate by the household's own transfer receipt.

Sixth, the design contains its own placebos, since salaried and retired households sit in the same districts and months as the exposed occupations and have no income channel through which rainfall can act, so every result for the exposed groups can be set beside a result for them that should be null, and Appendix Table A5 does this for the four main margins.

### 2.4 Credit in the Coping Portfolio

A weather shock reaches a household through what it does for a living, and the two tails of the distribution reach different occupations. A rainfall deficit works through the crop and through the wage that clears the local labour market when the crop suffers, so it lands on agricultural labourers and small farmers (Jayachandran, 2006), and where the wage cannot fall the adjustment moves into employment, so the deficit reaches landless labourers through fewer days as well as lower pay (Kaur, 2019). A rainfall surplus works through stopped outdoor work, so it lands on construction labourers and on the self-employed whose trade is outdoors, a channel that the crop-centred literature has had less occasion to test and that the earnings regressions of Section 4.1 test directly. Households whose income arrives as a salary or a pension are reached

by neither, which is what makes them placebos.

The household's first response is to change what it buys. The demand literature has long distinguished committed from deferrable expenditure (Deaton, 1997), Banerjee and Duflo (2007) document how much of a poor household's budget is already committed to food, fuel and rent, and Browning and Crossley (2009) show that a household facing a temporary loss of income smooths its welfare cheaply by postponing the purchase of durables and semi-durables while spending on what is consumed day to day holds. A transitory shock should therefore act on the deferrable margin first, and the cuts within the basket should run in the order of postponability, inside a total that barely moves, which is the pattern Malaeb et al. (2026) find on rural Indian panel data.

What a formal lender adds is the option to close the gap without going further down that order. In the buffer-stock model of Deaton (1991) a household that cannot borrow holds assets against bad draws and cuts consumption when they run out, and access to credit relaxes that constraint in exactly the states in which it binds, which is what the experiments that relax it find (Lane, 2024; Suri et al., 2021; Jack and Suri, 2014). The option need not be the household's own loan. When the lenders of a district lose their funding, the merchants who sell on account, the employers who advance wages and the small lenders that borrow from the same market lose working capital too, so a household that never held a non-bank loan can lose the slack it drew on after a bad month, which is the equilibrium channel Breza and Kinnan (2021) document for the microfinance contraction. The predictions that follow are therefore for the households of a district whose formal credit contracted, through their own lender and through the others. When formal credit contracts, the reallocation should deepen and, what matters more, extend into categories the household otherwise protects, and we name three in advance. Health and education are the expenses that Suri et al. (2021) find households with access stop forgoing, and the ones whose cuts the literature on rainfall shocks ties to lasting losses in children's schooling and health (Jensen, 2000; Maccini and Yang, 2009; Shah and Steinberg, 2017). Milk is the staple bought daily by nearly every household in our data. Together the three form the protected index that Section 3.2 names as a primary outcome.

The borrowing margins should move at the same time, and the sources should not move alike. Professional moneylenders lend their own capital and can take on displaced borrowers at a price, so that margin should rise where exposure was higher. Borrowing from relatives is the reciprocal insurance of Udry (1994) and Townsend (1994), which pools idiosyncratic risk inside a village and is poorly placed to insure a dry shock that hits every household in it, so that margin should answer the shock itself and not the credit regime. Many displaced households should stop borrowing altogether, because the informal sector cannot absorb a covariate shock and because Breza and Kinnan (2021) find no sign that informal lending filled the gap the microfinance contraction opened. Since the shock was to non-bank funding, bank borrowing should not fall. The margins beyond the budget may move in the same direction, since households have

sold bullocks when credit failed them (Rosenzweig and Wolpin, 1993), drawn down livestock (Kazianga and Udry, 2006) and supplied more labour as consumption fell (Jayachandran, 2006), so we also look at borrowing for medical bills and to service existing debt, the sale of assets, saving in gold, and the addition of a working member, expecting each to be imprecisely measured on its own. The order matters because its far end may not reverse, which is the asset-based poverty trap of Carter and Barrett (2006), and Section 4 reports what the panel can say about persistence within its own span.

One prediction precedes the rest, because the default did not change the weather or the occupational mix of any district, so the shock itself should be the same size in both regimes. The income loss of the exposed occupations after a dry shock, and the response of total spending, should not differ between districts with higher and lower exposure. If they did, the differences that follow could be a larger shock rather than a different way of absorbing the same one, as in the equilibrium channel of Breza and Kinnan (2021), and the comparison in Section 4 begins there. The account then yields seven further predictions, and Section 4 reports each beside its exhibit. Household earnings should fall after a weather shock by occupation and by tail, and not for salaried or retired households, which is the first. Total expenditure should not respond, which is the second, while the composition of the basket should, with the cuts running from the most postponable categories to the least, which is the third. Where exposure was higher, the cuts should deepen and extend into health, education and a staple such as milk, the fourth. Non-bank borrowing should fall and moneylender borrowing rise while bank borrowing holds, the fifth, and of the households that lose formal access, a large share should stop borrowing rather than switch to an informal source, the sixth. The seventh is the set of margins beyond the budget, which we expect to move together in the direction the account implies and to be individually imprecise.

# 3 Data and Empirical Strategy

## 3.1 Data

*Credit Bureau Records.* The credit data are a proprietary extract from TransUnion CIBIL, India's largest credit bureau, covering the formal retail lending that reaches the bureau, and close to a universe of it, since the extract's 2016 count of loans outside mortgages is within 3 percent of the count that Cramer et al. (2025) report from the same bureau. The records are cells of originations at pincode by month by product by lender type by credit-score band, with a flag for a sanctioned amount of fifty thousand rupees or more, aggregated to Census-2011 districts with a population-weighted pincode crosswalk. They span 640 districts from January 2015 to December 2019 and contain 483 million originations worth 105 trillion rupees. Lenders are classified as public banks, private banks, non-bank and housing-finance companies, financial-technology non-banks, and others, and products into thirteen categories. Two series are left

out of every total because the bureau's own reporting of them breaks inside the window, gold loans by non-bank lenders and the farm credit cards reported by lenders outside the four main groups, and 12 districts in which more than a twentieth of credit inquiries fall in pincodes that never record a loan are dropped from the lending regressions, since their placement is suspect. Six of them are in the household panel and keep their exposure shares, which are built from originations and not from inquiries. Appendix C gives the treatment of missing and stand-in amounts, re-reported loans and the lender taxonomy.

*Measuring Exposure.* For each district $d$ we compute the share of loans originated by non-bank and housing-finance lenders over the twelve months before the first sign of trouble,

$$\text{NBFC}_d = \frac{\sum_{m \in \text{June 2017 to May 2018}} \text{Loans}_{dm}^{\text{NBFC}}}{\sum_{m \in \text{June 2017 to May 2018}} \text{Loans}_{dm}}, \tag{1}$$

and we hold the share fixed for the rest of the analysis. A district in which, over those months, banks originated five of every ten retail loans and non-bank and housing-finance lenders the other five has a share of 0.50, which places it in the top tercile, and it keeps that value in every block of the panel, including the blocks before the default, when its lenders were still lending. Measuring exposure on the eve of the shock and holding it fixed is the practice of Breza and Kinnan (2021), who take each district's exposure to the Andhra Pradesh lenders in the quarter before the ordinance, and of Khwaja and Mian (2008) and Huber (2018), who take a borrower's lender relationships before the liquidity shock reached the lenders. The share has to be frozen because weather shocks move district credit, so a contemporaneous measure would condition on a variable the treatment moves. The window ends in May 2018, the month before the first default inside the group, so it predates any anticipation. The choice has one cost, since the window overlaps the last four pre-default blocks of the panel, and a dry month raises non-bank lending by about 3 percent (Table 2), so the share can embody the weather of 2017 and early 2018. A share measured over 2015 and 2016, which overlaps none of the last four pre-default blocks and predates the expansion of non-bank lending in 2017 and 2018 that the vehicle and property slowdown of 2019 then reversed, correlates 0.89 with it, places three districts in four in the same tercile, gives the same estimates in Section 4.3, and leaves the event study of Section 3.3 where it is (Appendix B). Two further exposures answer a different objection, that a high share mostly means many consumer-durable and vehicle loans, which are urban products. A share built within product and averaged at national product weights, and a share built on loans below fifty thousand rupees alone, the loans the exposed occupations hold, give the same estimates too.

Across the 640 districts the share has a mean of 0.34 and a standard deviation of 0.15, runs from 0.11 at the tenth percentile to 0.54 at the ninetieth, and reaches 0.72 at most (Appendix Figure A3). Districts are ranked on it, with tercile cutoffs at 0.28 and 0.41, and the main specifications compare the top tercile, whose mean share is 0.51, with the bottom tercile, whose

mean is 0.16, with the middle tercile carrying its own interactions. Appendix Table A3 compares the terciles on the districts that enter the estimates, and Section 2.3 read it. The household panel covers 518 of the districts with both an exposure share and a weather index, and it over-represents the exposed ones, with 206 of the 214 top-tercile districts, and 129 of the 213 bottom-tercile districts, of which 122 have a weather index, so the control group is the surveyed three fifths of the bottom tercile, which are larger than the unsurveyed two fifths. 10 of the 27 states in the panel hold districts in both terciles.

*The Household Panel.* The CMIE Consumer Pyramids Household Survey is a rotating panel of Indian households interviewed three times a year. Each interview collects consumption for the four preceding calendar months, income by source for the same months, and a set of wave-level modules on borrowing, assets, savings and the employment of every member. The unit of observation throughout is the interview block, the four reference months that one interview reports on. Our sample runs from January 2014 to December 2019 and covers 210,560 households in 525 Census-2011 districts, with 2,476,011 interview blocks, of which 2,313,807 report all four months. The estimation sample keeps the complete blocks in the 518 districts that have an exposure share and a weather index, and drops the blocks whose four months straddle the default, which leaves 209,874 households and 2,196,418 household-blocks. The regressions are unweighted, since their estimands are changes within households, and descriptive shares are unweighted sample shares.

Appendix Table A1 describes the sample by exposure group. The ten groups are built from the occupation of the household head at first observation, crossed with whether the head's industry puts the work outdoors, so that wage labourers in construction, mining, brick making, fishing, forestry and cultivation form the construction labourer group, the outdoor self-employed are those in building, animal husbandry, fishing and forestry, and traders, shopkeepers and hawkers fall with the other self-employed. The assignment is held fixed for the life of the household, so a household cannot move between treatment and control in response to the shock under study. The six groups with an income response in Section 4, agricultural labourers, small and organised farmers, construction and indoor casual labourers, and the outdoor self-employed, are the exposed occupations, with 83,263 households in the panel and 82,992 in the estimation sample, and salaried and retired households are the two placebos. Income is recorded by source and is smoothed in reporting, with a median month-on-month change within a block of 3 percent against 11 percent for spending, so the income effect below is attenuated by construction and is a lower bound on the loss.

Consumption is recorded in far more detail than borrowing. The survey records over a hundred and fifty individual spending items, monthly and in rupees, fifteen heads partition the total, and from these we take twenty categories that cover 50.7 plus 26.7 percent of the budget, an essential set and a discretionary set, with the remaining 22.6 percent, other food, loan instalments and the miscellaneous head, carried as a residual. Appendix Table A4 lists

the categories with their shares, the share of household-months in which each is bought, and their Engel elasticities. The division into essential and discretionary follows the criterion of postponability of Section 2 and was fixed before estimation. It is not an income elasticity, and the two disagree for eight of the twenty categories, since intoxicants, cosmetics, transport and snacks behave like necessities in the cross section while milk, health, power and bills behave like luxuries, and where the two concepts diverge we go with deferability, because a transitory shock acts on that margin. The classification can be argued with, so we report the twenty categories separately throughout. The survey records rupees and not quantities, so a flat rupee figure for cereals after a dry shock that raises cereal prices is consistent with fewer kilograms, and Section 4 reports the extensive margin, whether a household bought a category at all, where the distinction matters.

Borrowing is a matrix of binary indicators crossing eleven sources with eleven purposes, recorded once per wave as a stock on the date of the interview. The survey records no amounts, interest rates or maturities, so every credit result in this paper is an extensive-margin result about outstanding borrowing, and the paper will speak of switching and never of substitution. The module was built up over the panel, so every borrowing result uses waves 13 to 18, January 2018 to December 2019, which leaves two waves before the default. From wave 15, the wave that contains the default, the questionnaire added a box for a loan whose source the household cannot name, and we leave that box out of every source and state, report the borrowing estimates without wave 15 in Appendix B, and show in Appendix Figure A5 that no named source jumps at the wave in either tercile. Borrowing from shops is recorded only when a payment to a shop is overdue by two months or more, so it is overdue trade credit and a distress indicator in its own right. The survey carries one rupee measure of debt, the instalments a household pays each month on its loans, reported by 9.3 percent of household-months with a median of 1,800 rupees, in total and for housing, vehicle and consumer-durable loans separately. Income from the sale of assets is recorded for 2.6 percent of household-waves in 2018 and 2019 and for almost none before, so it is estimated on the same waves as borrowing, saving in gold is a flow indicator for the preceding four months, and the added-worker margin is a rise in the household's count of members in work between waves, available from 2016. Households report borrowing from an entity described as a non-bank or dealer, a single category that covers a finance company, a fintech lender and a shop selling on instalment alike, so claims about which kind of lender withdrew rest on the bureau. The panel under-represents the poor (Somanchi, 2021; Drèze and Somanchi, 2024), so every estimand in the paper is a change within a household or a difference between groups, which survive non-differential bias (Abraham and Shrivastava, 2022), and the magnitudes apply to the surveyed population. Appendix C gives the remaining detail on the modules, the sampling frame and its 2017 expansion, and the district crosswalk.

*The Weather Index.* We measure weather with the Standardised Precipitation and Evapotranspiration Index of Vicente-Serrano et al. (2010), from SPEIbase v2.9 at half-degree resolution

and monthly frequency, aggregated to districts with 2011 gridded population weights. The index nets potential evapotranspiration out of precipitation and standardises the result against the distribution for that cell and calendar month, so that zero is a normal month for that place at that time of year. The treatment is a pair of indicators,

$$\text{Dry}_{dm} = \mathbf{1}[\text{SPEI}_{dm} < q_d(0.20)]\,, \qquad \text{Wet}_{dm} = \mathbf{1}[\text{SPEI}_{dm} > q_d(0.80)] \tag{2}$$

where $q_d(p)$ is the $p$th quantile of district $d$'s index over January 2001 to December 2021, the window Cramer et al. (2025) use, following the practice of defining shocks against local history (Jayachandran, 2006; Corno et al., 2020). Over 2014 to 2019 the dry tail fires in 19.0 percent of district-months and the wet tail in 20.4 percent, so the study years are not unusual ones, and thresholds over the full record from 1901 are a robustness check. 18 of the 641 Census district codes have no index, 17 of which are bureau districts and 14 of which lie in Jammu and Kashmir, because SPEIbase masks the disputed territory. Because a block spans four months, a block is treated if the dry tail fires in any of them, which happens in 50 percent of district-blocks, so the treatment is common, a treated block holds one and a half dry months on average, and we call such a block a dry shock, which is a dry spell inside the block and not a failed season, with the seasonal windows of Appendix B as the check that a longer deficit gives the same ordering. Appendix Table A2 counts the district-blocks in each cell of the design. The two tails are not equally well measured, since a rank test against documented events, reported in Appendix D (Table A14), places the 2002 and 2009 droughts third and eighth among 243 months while four major floods rank between thirty-sixth and forty-eighth for their own regions, so we interpret the wet tail as an unusually wet month and not as a flood, and we keep the two tails apart rather than pooling them. Appendix Figures A1 and A2 map the exposure terciles and the post-default dry-tail frequency across districts and show when each tail fired. Shocks arrive in national clusters, and the state-by-block effects absorb months in which nearly every district in a state is in the same tail, so the identifying variation comes from the months in which districts within a state differ.

### 3.2 Empirical Strategy

We estimate at the frequency of the interview block because of a property of the survey instrument. A household is interviewed once and asked to recall the four preceding months, so four monthly rows come from a single interview, between 4 and 16 percent of blocks report the same spending in all four months, and the demeaned means by recall position alternate instead of decaying. Treating the rows as independent readings would overstate the sample fourfold while mistaking recall structure for economic timing, and the alternating pattern cancels exactly in equal-weighted complete-block means. The estimation sample for income and spending is therefore 2,196,418 household-blocks, the borrowing, asset-sale and gold margins use the 880,588

household-waves of waves 13 to 18, and the added-worker margin the 1,510,197 household-waves of waves 8 to 18, each wave compared with the one before. The post-default indicator is one for blocks whose reference months all fall in September 2018 or later and zero for blocks whose months all fall before, a block is assigned to the calendar block, and so to the state-by-block effect and the event-study index, that holds three or more of its four reference months, and a block split two and two is assigned to the calendar block in which the interview falls. The wave-level outcomes, borrowing, asset sales, gold saving and the added member, are stocks recorded on the date of the interview, so for them the post-default indicator is one for interviews after 4 September 2018, waves 15 to 18, the event-study index is the interview wave, and the weather shock is the dry tail over the four reference months the interview reports on. Category spending enters as the inverse hyperbolic sine of rupees, which keeps the household-blocks with no purchase, and its coefficients are read as approximate percent changes. Chen and Roth (2024) show that such a reading is not invariant to the unit of the outcome when the treatment moves households between zero and positive spending, so for the protected categories, where more than two fifths of the household-blocks record no education spending, the paper reports the extensive margin separately and, in Appendix B, the same rows by Poisson pseudo-maximum likelihood, which is unit-free. In November 2016 the government withdrew the two largest banknotes and every group's reported income falls by 9 to 14 percent, salaried and retired households included, which the state-by-block effects absorb.

The earnings regression is

$$\ln Y_{hdt} = \beta^{D}\,\mathrm{Dry}_{dt} + \beta^{W}\,\mathrm{Wet}_{dt} + \alpha_h + \tau_{st} + \varepsilon_{hdt} \tag{3}$$

where $h$ indexes households, $d$ districts, $t$ four-month blocks and $s$ states. This is the specification with which Jayachandran (2006) and Shah and Steinberg (2017) read district rainfall shocks into wages and household outcomes, a shock defined in the tails of the district's own rainfall history with unit and time effects, except that their shock is a year of rainfall and ours is a month of the water balance, the frequency at which Cramer et al. (2025) measure the lending side of the same episode, with the household fixed effect $\alpha_h$ in place of the district effect because the panel follows households, so nothing in the results can be driven by caste, landholding, education or distance to a bank. The state-by-block effect $\tau_{st}$ removes inflation, state policy and the national cycle. We estimate equation (3) separately by exposure group, and the placebo comparison sits inside this split.

For each margin of the portfolio we estimate the same equation with the margin in place of income, then interact the shock with the credit contraction,

$$\begin{aligned} Y_{hdt} = {} & \beta_s\,S_{dt} + \theta\left(S_{dt}\times \mathrm{NBFC}_d \times \mathrm{Post}_t\right) + \lambda\left(S_{dt}\times \mathrm{NBFC}_d\right) \\ & + \mu\left(\mathrm{NBFC}_d \times \mathrm{Post}_t\right) + \rho_s\left(S_{dt}\times \mathrm{Post}_t\right) + \alpha_h + \tau_{st} + \varepsilon_{hdt} \end{aligned} \tag{4}$$

where $S_{dt}$ is the weather shock, $\text{NBFC}_d$ is the frozen reliance of equation (1), entered as the top-tercile indicator with the middle tercile carrying its own set of interactions, $\text{Post}_t$ indicates blocks from September 2018 onward, and $\beta_s$ and $\rho_s$ carry a state subscript, so that the weather response and its change at the default are allowed to differ across states. Both tails enter every regression with their own interactions, and the tables report the dry-tail terms. The specification joins two designs from the literature, since the interaction of a shock with access to finance is the design of Jack and Suri (2014) and Suri et al. (2021), who ask whether a household with access cuts consumption less after a shock than one without, and the interaction of frozen district exposure with a post-event indicator is the design of Breza and Kinnan (2021) for the microfinance contraction. Cramer et al. (2025) combine the two on the lending side of the same episode, and equation (4) is the household counterpart of their specification. The level of $\text{NBFC}_d$ is absorbed by the household effect and the level of $\text{Post}_t$ by the state-by-block effect. The term $\lambda$ absorbs any tendency of non-bank-reliant districts to respond differently to weather for reasons unrelated to the default, $\mu$ absorbs any post-default divergence in those districts unrelated to weather, and $\rho_s$ absorbs any change after the default in the weather response that is common to the districts of a state. The coefficient $\theta$ is therefore identified from the triple difference within states, from the 10 of the 27 states that hold districts in both terciles, and it measures the change at the default in the effect of the same weather shock, on the same margin, in a district with higher exposure to the contraction relative to one in the same state with lower exposure. Section 4.3 reports the pooled design in which $\beta$ and $\rho$ are common across states, so that the cross-state variation enters as well, and it gives an estimate a quarter larger. The identifying assumption is that households in districts with the same weather history and the same occupational mix would have changed their response to weather at the default in the same way regardless of whether the lenders in the district were non-banks. Section 3.3 offers the evidence we have for it, and Section 4.3 the checks that would detect its failure.

The tables that follow report three quantities for each margin. The column headed lower exposure is $\hat{\beta} + \hat{\rho}$, the response to a dry shock after the default in the bottom tercile, where $\hat{\beta}$ and $\hat{\rho}$ are the averages of the state-specific coefficients weighted by district-blocks. The column headed higher exposure is $\hat{\beta} + \hat{\rho} + \hat{\theta}$, the response the top tercile shows once its pre-default differential $\hat{\lambda}$ is removed, so that the difference between the two columns is $\hat{\theta}$ with its standard error, the standard errors of the sums come from the estimated covariance matrix, and Appendix Table A7 reports all five coefficients for the main margins. The labels are the two ends of a contrast between a smaller and a larger contraction, since the bottom tercile also lost non-bank lending after the default, by 12 log points relative to banks in the same districts, with a standard error of 9 that admits no loss, against 27 log points, about 24 percent, in the top tercile (Table 1, Panel B), and $\hat{\mu}$, the change in the outcome in the top tercile after the default in a block without a shock, is the survey's own reading of the contraction. Three outcomes were named as primary in a design note written before estimation, the discretionary budget

share, the share of households holding a non-bank loan, and the protected index, the log of rupees spent on health, education and milk together, and the portfolio tables report beside each $\theta$ the sharpened false discovery rate $q$-value of Anderson (2008) computed across the twenty rows of Tables 6 and 7 together. Standard errors are clustered by district in every table, as in Breza and Kinnan (2021), and Appendix Table A11 reports the main estimate under spatial standard errors in the manner of Conley (1999), clustering by the survey's homogeneous regions and by state with the wild cluster bootstrap of Cameron et al. (2008), and a permutation test that reassigns the exposure terciles across districts within state, as Breza and Kinnan (2021) permute exposure across districts. The weather treatment switches on and off, so we also report the de Chaisemartin and D'Haultfœuille (2020) estimator for switching treatments, applied to district-block means within each cell of the design with $\theta$ as the double difference, as a check on the two-way fixed-effects specification, with the construction in Appendix B.

### 3.3 Pre-Trend Analysis

If the default was an exogenous shock to credit supply in exposed districts, the weather response of exposed and unexposed districts should not have diverged before it. Figure 2 estimates

$$\begin{aligned} Y_{hdt} = & \sum_{k \neq -1} \delta_k \left( S_{dt} \times \mathrm{NBFC}_d \times \mathbf{1}[t = k] \right) + \sum_{k \neq -1} \gamma_k \left( \mathrm{NBFC}_d \times \mathbf{1}[t = k] \right) \\ & + \sum_{k \neq -1} \pi_{sk} \left( S_{dt} \times \mathbf{1}[t = k] \right) + \beta_s \, S_{dt} + \lambda \left( S_{dt} \times \mathrm{NBFC}_d \right) + \alpha_h + \tau_{st} + \varepsilon_{hdt} \end{aligned} \tag{5}$$

where $k$ indexes four-month blocks of reference months relative to the block September to December 2018, and $k = -1$, May to August 2018, is the omitted reference, in the event-study form in which Breza and Kinnan (2021) display exposure interacted with survey round. The $\gamma_k$ terms allow non-bank-reliant districts to follow their own path independently of the weather, and the $\pi_{sk}$ terms, which differ by state as $\rho_s$ does in equation (4), allow the weather response of the districts of each state to follow its own path independently of exposure. Block $k = 0$ is fully treated for the block outcomes, since its reference months all fall after the default, and for the wave-level outcomes, since the interviews of wave 15 begin after it, though the funding run developed over the following three months, so the block is a transition in the lending data. The figure tracks six margins, log income for the exposed occupations, the discretionary share, the protected index and instalments paid from January 2016, eight blocks before the block that contains the default, with the blocks of 2014 and 2015 binned into a single endpoint indicator so that the pre-period stays readable and $k = -1$ remains the only omitted block, and non-bank and moneylender borrowing over the six interview waves in which the borrowing module is used, indexed by interview wave, which leaves those two panels a single pre-period coefficient.

For the discretionary share, the seven pre-period coefficients are small and individually indistinguishable from zero, the largest of them, three blocks before the default, at 0.45 with a

standard error of 0.39. A joint test of the seven pre-period coefficients does not reject, with a $p$-value of 0.27, dropping the block with the largest coefficient leaves $\theta$ at $-0.50$ with a standard error of 0.32 (Appendix Table A10), while the sensitivity analysis of Rambachan and Roth (2023) adds nothing here. Under its relative-magnitudes restriction the post-default differential may drift each block by a multiple of the largest change between consecutive pre-period coefficients, and that change is 0.55 points, into the block with the largest coefficient, which exceeds the effect, while the headline's 95 percent interval already reaches zero before any drift is allowed, so the pre-period coefficients, with standard errors of four tenths of a point each, are too imprecise to bound the drift, and the evidence for the identifying assumption is the joint test, the stability of $\theta$ to dropping the block, and the fact that the main estimates are identified within states, and no more than that. For log income of the exposed occupations the coefficients are flat before and after the default, with a joint pre-period $p$-value of 0.90, which is the first look at the prediction that the shock itself did not change size, and for the protected index the pre-period is flat, the point estimates open from the first post-default block, and only the last is distinguishable from zero at five percent. After the default the discretionary gap is $-0.33$ points in the block that contains the default, $-0.56$ in the next, $-0.50$ in the block after that and $-0.69$ in the last block the panel allows, September to December 2019, and the four average to the $\theta$ of Table 5. The three later coefficients are not distinguishable from one another, and the dip falls in May to August 2019, the half-year in which the funding squeeze was most acute by the account of Sengupta et al. (2022), so it runs against the funding chronology and we have no explanation for it. Instalments paid are noisy before the default and fall after it, reaching 10 percent in the last block. The point estimates for non-bank borrowing fall from the first post-default wave and moneylender borrowing rises in step, with no single wave distinguishable from zero at the precision one wave affords, the last two non-bank coefficients being $-0.77$ and $-0.88$ with standard errors of 0.63 and 0.67, which is why the estimates of Table 6, which average over the post-default waves, carry the result, and their one pre-period coefficient each is small. The figure also shows the main result in its simplest form, since the gap that opens after the default is the extra cut in discretionary spending that a dry shock produces where exposure was higher.

# 4 Results

## 4.1 The Contraction, the Weather and Earnings

The event study already shows the main result in qualitative form. We turn to the three relationships the design rests on, from the funding shock to district lending, from the weather to district lending, and from the weather to household earnings.

*The Contraction in Exposed Districts.* Table 1 gives the effect of the default on district lending, which scales the household estimates, from a regression of district lending by month on exposure interacted with the post-default indicator, with district and month fixed effects, which is the

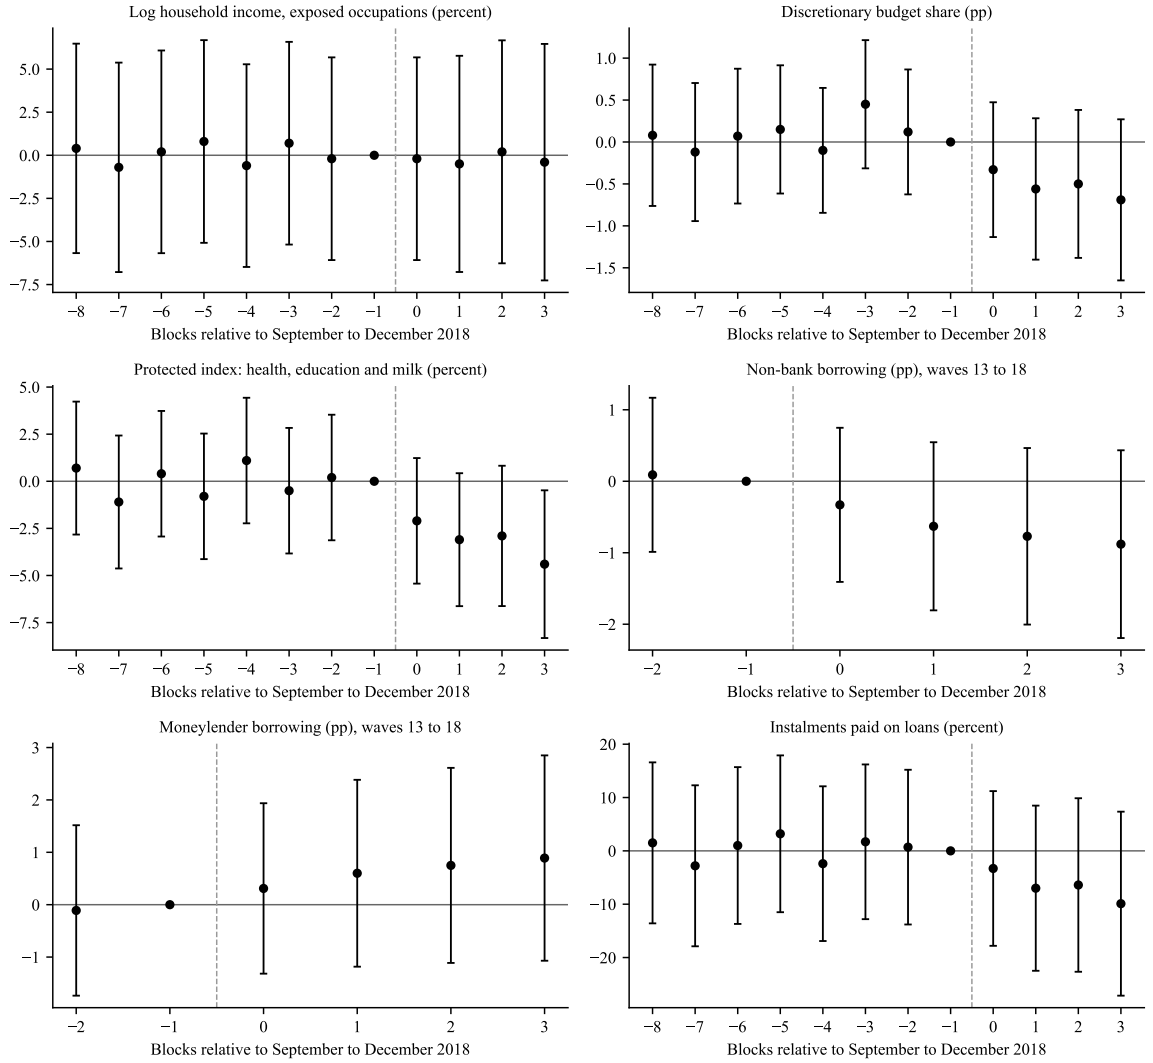



**Figure 2:** Event Study around the September 2018 Default

Coefficients $\delta_k$ from equation (5) with 95 percent intervals, standard errors clustered by district. The omitted block is May to August 2018, and the dashed line marks the default. The income panel uses the exposed occupations. The income, discretionary-share, protected-index and instalment panels run from January 2016, with the blocks of 2014 and 2015 binned into one endpoint indicator that is not shown, and the borrowing panels use interview waves 13 to 18, so their first coefficient is the wave interviewed in January to April 2018. Joint tests of the pre-period coefficients give $p$-values of 0.90 for income, 0.27 for the discretionary share and 0.70 for the protected index.

lending-side specification of Breza and Kinnan (2021).

The first row of Panel A is the estimate proper. Non-bank lenders in the top tercile of pre-default reliance lent 12 percent fewer rupees after the default relative to the bottom tercile, significant at one percent, and a one standard deviation higher share lowers their lending by 6 percent (column 2). The fall is the same for loans above and below fifty thousand rupees and the same in counts as in rupees, so non-banks made fewer loans of the same size. Their share of the district's rupees falls by 4.4 percentage points more, from 33.4 percent in July 2018 to

**Table 1:** The Contraction in Exposed Districts

| Outcome, Panel A | (1) Top tercile × Post | (2) Share, 1 s.d. × Post | (3) Top tercile, district trends | (4) Placebo, Post from September 2017 |
|---|---|---|---|---|
| Non-bank rupees (log) | -0.132*** | -0.059** | -0.071 | -0.033 |
| | (0.049) | (0.025) | (0.060) | (0.052) |
| Non-bank loans (log) | -0.130** | -0.058** | -0.066 | -0.041 |
| | (0.056) | (0.028) | (0.066) | (0.058) |
| of which below Rs 50,000 (log) | -0.138** | -0.061** | -0.072 | -0.045 |
| | (0.062) | (0.031) | (0.072) | (0.064) |
| of which Rs 50,000 or more (log) | -0.121** | -0.054* | -0.061 | -0.036 |
| | (0.061) | (0.031) | (0.071) | (0.063) |
| Non-bank share of rupees (pp) | -4.4*** | -2.0*** | -2.5* | -0.8 |
| | (1.2) | (0.6) | (1.4) | (1.3) |
| All rupees (log) | -0.074* | -0.033* | -0.041 | -0.009 |
| | (0.038) | (0.019) | (0.045) | (0.040) |
| All loans (log) | 0.124*** | 0.055** | 0.082 | 0.027 |
| | (0.047) | (0.024) | (0.054) | (0.049) |
| Bank rupees (log) | -0.038 | -0.017 | -0.021 | 0.006 |
| | (0.044) | (0.022) | (0.049) | (0.046) |
| Bank loans (log) | 0.088* | 0.039 | 0.058 | 0.021 |
| | (0.049) | (0.025) | (0.056) | (0.051) |
| Bank rupees per loan (log) | -0.126** | -0.056** | -0.079 | -0.015 |
| | (0.052) | (0.026) | (0.060) | (0.054) |
| Fintech loans (inverse hyperbolic sine) | 0.312*** | 0.139** | 0.244* | 0.088 |
| | (0.121) | (0.061) | (0.138) | (0.127) |
| *Panel B. Within district, lender type × Post, banks omitted (log rupees)* | | | | |
| NBFC/HFC × Post, district by month effects | -0.21*** | | | |
| | (0.06) | | | |
| NBFC/HFC × Post, district by product by month effects | -0.18*** | | | |
| | (0.06) | | | |
| Fintech × Post | 0.29*** | | | |
| | (0.11) | | | |
| Other lenders × Post | -0.08 | | | |
| | (0.09) | | | |
| NBFC/HFC × Post, top-tercile districts | -0.27*** | | | |
| | (0.08) | | | |
| NBFC/HFC × Post, bottom-tercile districts | -0.12 | | | |
| | (0.09) | | | |
| District-months, Panel A | 37,680 (628 districts × 60 months) | | | |

Panel A: each cell is a separate regression of the row outcome, at district by month, on the exposure measure interacted with an indicator for September 2018 onward, with district and month fixed effects, over January 2015 to December 2019. Log outcomes are of rupees or loans originated in the district-month, the share is in percentage points, and the fintech row uses the inverse hyperbolic sine because most district-months before 2017 record no fintech loan. Column (1) interacts the top-tercile indicator with the post-default indicator, column (2) the share standardised to unit variance, column (3) adds district-specific linear trends to column (1), and column (4) moves the post-default date to September 2017 and ends the sample in August 2018. Panel B regresses log rupees originated at district by lender type by month on lender-type indicators interacted with the post-default indicator, with lender-type effects and district-by-month effects in the first row and district-by-product-by-month effects in the second, banks omitted, in the manner of Khwaja and Mian (2008). The twelve districts whose inquiries cannot be placed by pincode are excluded. Standard errors clustered by district in parentheses. *, ** and *** denote significance at the 10, 5 and 1 percent levels.

27.0 three months later in the top tercile against 19.1 to 17.2 in the bottom. Total rupees lent in exposed districts fall by 7 percent relative to unexposed ones, and bank rupees by 4 percent with a standard error that admits no fall and a fall twice that size, so banks did not make up the difference in rupees. The count rows record the first result that does not fit a simple contraction. The number of loans in exposed districts rose by 13 percent relative to unexposed ones after the

default, bank loan counts by 9 percent, and the counts of financial-technology lenders by about a third, while bank rupees per loan fell by 12 percent, so the exposed districts ended the period with more loans and fewer rupees. The lenders whose funding did not depend on the commercial paper market filled part of the gap with small loans, which is also what Cramer et al. (2025) find. Panel B shows the same contraction within districts, where the non-bank group lent a fifth less than banks in the same district and month after the default, a sixth less within the same product, and more in the top tercile than in the bottom. The third column is the caveat of Section 2.3 in numbers. Allowing each district its own linear trend roughly halves the non-bank estimates and leaves them imprecise, because non-bank lending in low-reliance districts had been growing faster before the default, and the placebo in the fourth column, which dates the default a year early, is null throughout. The lending effect is therefore a break in the level of non-bank lending in exposed districts superimposed on a relative drift in the same direction. The household estimates below do not rest on the level of lending but on the change in the weather response at the default, which a drift in lending does not touch directly.

*Weather Shocks and District Lending.* The premise of the paper is that formal credit answered dry months before the default, and Table 2 tests it in the bureau with the lending-side counterpart of equation (4), the specification of Cramer et al. (2025), with district exposure in place of their lender type.

**Table 2:** Weather Shocks and District Lending, before and after the Default

| Outcome, district by month | (1) Dry shock, before the default | (2) Dry × top tercile, before the default | (3) Dry × top tercile × Post | (4) Dry × Post, all districts | (5) Post-default response, top tercile |
|---|---|---|---|---|---|
| Non-bank loans (log) | 0.031** | 0.006 | -0.046** | 0.004 | -0.005 |
| | (0.014) | (0.019) | (0.021) | (0.016) | (0.020) |
| of which below Rs 50,000 (log) | 0.044** | 0.009 | -0.058** | 0.006 | 0.001 |
| | (0.018) | (0.023) | (0.026) | (0.021) | (0.025) |
| Non-bank rupees (log) | 0.019 | 0.004 | -0.036 | 0.002 | -0.011 |
| | (0.016) | (0.021) | (0.024) | (0.018) | (0.023) |
| Bank loans (log) | 0.012 | -0.003 | -0.008 | 0.003 | 0.004 |
| | (0.011) | (0.015) | (0.016) | (0.013) | (0.016) |
| Bank rupees (log) | 0.008 | 0.001 | -0.005 | -0.001 | 0.003 |
| | (0.013) | (0.017) | (0.019) | (0.015) | (0.018) |
| Fintech loans (inverse hyperbolic sine) | 0.052* | 0.011 | 0.014 | 0.009 | 0.086* |
| | (0.031) | (0.040) | (0.046) | (0.036) | (0.045) |
| District-months | 36,660 (the 611 districts of the lending regressions that have a weather index × 60 months) | | | | |

Each row is a separate regression of the outcome, at district by month, on the dry-tail indicator, its interaction with the top-tercile indicator, and the triple interaction with the post-default indicator, with every lower-order term, district fixed effects and month fixed effects, over January 2015 to December 2019 in the districts with a weather index. Column (1) is the response of originations to a dry month before the default outside the top tercile, column (2) the pre-default difference in that response in the top tercile, column (3) the change in that difference after the default, column (4) the change after the default in the response outside the top tercile, and column (5) the post-default response in the top tercile, the sum of the four, with its standard error from the covariance matrix. Rows in logs drop the district-months with no origination in the row's series, under one percent of district-months in every row. Standard errors clustered by district in parentheses. *, ** and *** denote significance at the 10, 5 and 1 percent levels.

Before the default a dry month raised non-bank loan counts by 3 percent, and counts of non-bank loans under fifty thousand rupees by 4 percent, with the rupee response smaller and

imprecise, so the lending that answered a dry month was small-ticket lending, as it is in Cramer et al. (2025). Bank originations did not respond to dry months in either period. Exposed districts did not differ in the pre-default response, which is column (2), the response outside the top tercile did not change at the default, which is column (4), and the post-default response in exposed districts, the sum in column (5), is zero within a standard error of two percent, so that a dry month in an exposed district after September 2018 no longer drew the small-ticket non-bank lending it had drawn before. This is the supply-side version of the household result below, and also the ceiling on the direct channel, since the lending that stopped answering dry months was a few percent of a small base.

*Weather Shocks and Household Earnings.* Table 3 reports equation (3) by exposure group, and the pattern across groups, more than any single estimate, is the finding.

**Table 3:** Weather Shocks and Household Earnings, by Exposure Group

| | Log household income | |
|---|---|---|
| Exposure group | Dry tail | Wet tail |
| Agricultural labourer | -0.0621*** | -0.0374** |
| | (0.0183) | (0.0170) |
| Small farmer | -0.0483*** | -0.0261* |
| | (0.0142) | (0.0138) |
| Organised farmer | -0.0414** | -0.0108 |
| | (0.0196) | (0.0187) |
| Self-employed, outdoor | -0.0307 | -0.0479** |
| | (0.0214) | (0.0203) |
| Construction labourer | -0.0192* | -0.0552*** |
| | (0.0101) | (0.0129) |
| Casual labourer, indoor | -0.0238** | -0.0207* |
| | (0.0113) | (0.0109) |
| Self-employed, other | -0.0141 | -0.0119 |
| | (0.0088) | (0.0085) |
| Other | -0.0126 | -0.0098 |
| | (0.0129) | (0.0125) |
| *Salaried* | -0.0021 | 0.0027 |
| | (0.0068) | (0.0066) |
| *Retired* | 0.0009 | -0.0148* |
| | (0.0079) | (0.0076) |
| Household fixed effects | Yes | Yes |
| State × block fixed effects | Yes | Yes |
| Frequency | Four-month interview blocks | |
| Clustering | Census-2011 district | |
| Household-blocks, all groups | 2,196,418 | |

Estimates of $\beta^D$ and $\beta^W$ from equation (3), estimated separately by exposure group. Dependent variable is log household income, and the coefficients are read in the text as approximate percent changes. Both tails are entered in the same regression. Standard errors clustered by Census-2011 district in parentheses. *, ** and *** denote significance at the 10, 5 and 1 percent levels. Italicised groups are placebos.

A rainfall deficit lands on agricultural labourers, who lose 6.2 percent of household income,

and on small farmers, who lose 4.8 percent, with organised farmers at 4.1 percent. A rainfall surplus works instead through stopped outdoor work, and lands on construction labourers, who lose 5.5 percent, and on the outdoor self-employed, who lose 4.8 percent, while organised farmers barely move in the wet tail at all. This is prediction (1), and the asymmetry matters beyond this paper, since an indicator built from both tails of a standardised index, now the standard treatment variable in this literature, averages two events that reach different workers through different channels. Salaried households are indistinguishable from zero in both tails, as the design requires, while retired households are not, with a wet-tail estimate of $-0.0148$ and a standard error of 0.0076, marginally significant and out of place, since a retired household has no wage income for a stopped work site to remove. Appendix E discusses the possible reasons, none of which the data can test, and the result is the most serious challenge in the paper to the placebo logic. The retired placebo is null on income in the dry tail, which is the tail that carries the credit result, and null on every margin of the portfolio in Appendix Table A5.

### 4.2 The Coping Portfolio

*Total Expenditure and Budget Composition.* We turn to how much households cut, and Table 4 shows that on average they barely do, for the exposed occupations as for everyone. A dry shock lowers the income of the exposed occupations by 3.8 percent and their total spending by 0.7 percent, with the spending estimate indistinguishable from zero, and the all-household figures are 1.7 and 0.4 percent, smaller because three households in five have no income channel, and this is prediction (2). Meanwhile the discretionary share of the budget falls by 0.94 percentage points after a dry shock in the exposed occupations and by 0.68 in the all-household sample, while the residual share and the instalment share do not move. The shares sum to one whatever the total does, so a fall in the discretionary share is a rise in the other two by construction, and the content of the result is where the fall lands. A household that responds to a shock by buying different things instead of fewer things is invisible to a regression of total spending on rainfall, and a literature that runs that regression and finds nothing will conclude that households are well insured when they are reallocating. The total is also imprecisely measured, since a fall of one percent sits inside its interval, and the paper's claim is the narrower one that the composition moves by far more than the level.

*The Same Shock in Both Regimes.* Everything that follows compares the response to a dry shock where exposure to the contraction was higher with the response where it was lower, and the comparison means nothing unless the shock itself was the same size in both. Table 5 reports equation (4) on income and total spending, for the exposed occupations and for all households, in the three-column form of Section 3.2. A dry shock after the default lowers the income of the exposed occupations by 3.8 percent where exposure was lower and by 4.0 percent where it was higher, a difference of 0.2 percent with a standard error of 1.6, and wage income by 4.4 and 4.8 percent. Total spending moves by less than one percent in both regimes and the difference is

**Table 4:** Effects of Weather Shocks on Income, Total Expenditure and Budget Shares

| | Exposed occupations | | All households | |
|---|---|---|---|---|
| Outcome | Dry tail | Wet tail | Dry tail | Wet tail |
| Log household income | -0.0382*** | -0.0341*** | -0.0171*** | -0.0158*** |
| | (0.0097) | (0.0094) | (0.0062) | (0.0060) |
| Log total expenditure | -0.0071 | -0.0052 | -0.0043 | -0.0038 |
| | (0.0084) | (0.0081) | (0.0061) | (0.0058) |
| Discretionary share, pp | -0.94*** | -0.79*** | -0.681*** | -0.594*** |
| | (0.26) | (0.25) | (0.213) | (0.207) |
| Essential share, pp | 0.99*** | 0.83*** | 0.744*** | 0.627*** |
| | (0.27) | (0.26) | (0.221) | (0.216) |
| Residual share, pp | -0.05 | -0.04 | -0.063 | -0.033 |
| | (0.19) | (0.18) | (0.148) | (0.144) |
| Instalment (EMI) share, pp | -0.06 | -0.03 | -0.047 | -0.021 |
| | (0.07) | (0.07) | (0.052) | (0.051) |
| Baseline discretionary share | 26.3% | | 26.7% | |
| Baseline essential share | 50.7% | | 50.7% | |
| Households | 82,992 | | 209,874 | |
| Household-blocks | 863,296 | | 2,196,418 | |

Each cell is a separate estimate of $\beta^D$ or $\beta^W$ from equation (3) with the stated outcome in place of log income. The exposed occupations are agricultural labourers, small and organised farmers, construction and indoor casual labourers, and the outdoor self-employed. All specifications include household and state by block fixed effects at four-month frequency, with standard errors clustered by district. Shares are in percentage points of total monthly expenditure and sum to one hundred across the three share rows, so their coefficients sum to zero. The residual share is the budget outside the twenty categories, and the instalment share is the survey's equated monthly instalment head, which sits inside the residual.

0.2 percent with a standard error of 1.3, and the same holds for the all-household sample. The shock was the same size where exposure was higher, within the precision of these data, which is limited, since a standard error of 1.6 on a loss of 3.8 percent cannot rule out a shock close to twice as large in exposed districts, and the discretionary share nonetheless fell by 0.85 points more in the exposed occupations and by 0.52 points more overall. What the contraction changed was how the loss was absorbed, not its size, and the rest of this section is about the how.

*Category-Level Responses.* If the total holds still while the shares move, the question becomes which categories give way, and whether credit changes the list. Figure 3 reports the response of each of the twenty basket categories to a dry shock after the default, where exposure was lower and where it was higher, in Panel A, and the difference between the two, which is $\theta$, in Panel B. The categories are ordered by the all-household dry-tail response of Appendix Figure A4, which is prediction (3) and, read from the bottom of the figure upward, runs from the most postponable to the least. Recreation falls 14 percent in the all-household estimate, appliances 12, restaurant meals 12, gold and jewellery 11, with clothing and footwear behind them, and at the far end cereals and power do not fall. Because the order within the blocks was not fixed in advance, we check it against an order that was. Ranking the twenty categories by the durability

**Table 5:** The Same Dry Shock under Lower and Higher Exposure

| | Baseline | Effect of a dry shock after the default | | |
|---|---|---|---|---|
| | | Lower exposure | Higher exposure | Difference, $\theta$ |
| *Panel A. Exposed occupations (82,992 households, 863,296 household-blocks)* | | | | |
| Log household income (×100) | | -3.82*** | -4.02*** | -0.20 |
| | | (1.03) | (1.44) | (1.62) |
| Log wage income (×100) | | -4.41*** | -4.78*** | -0.37 |
| | | (1.32) | (1.84) | (2.06) |
| Log total expenditure (×100) | | -0.61 | -0.83 | -0.22 |
| | | (0.84) | (1.16) | (1.25) |
| Discretionary share, pp | 26.3 | -0.72*** | -1.57*** | -0.85** |
| | | (0.27) | (0.39) | (0.43) |
| Residual share, pp | 23.0 | -0.04 | 0.15 | 0.19 |
| | | (0.20) | (0.29) | (0.32) |
| Essential share, pp, implied | 50.7 | 0.76*** | 1.42*** | 0.66 |
| | | (0.28) | (0.41) | (0.45) |
| *Panel B. All households (209,874 households, 2,196,418 household-blocks)* | | | | |
| Log household income (×100) | | -1.66** | -1.83** | -0.17 |
| | | (0.65) | (0.93) | (1.05) |
| Log total expenditure (×100) | | -0.39 | -0.50 | -0.11 |
| | | (0.62) | (0.85) | (0.93) |
| Discretionary share, pp | 26.7 | -0.42** | -0.94*** | -0.52* |
| | | (0.19) | (0.27) | (0.31) |
| Residual share, pp | 22.6 | -0.05 | 0.02 | 0.07 |
| | | (0.15) | (0.20) | (0.23) |
| Essential share, pp, implied | 50.7 | 0.47** | 0.92*** | 0.45 |
| | | (0.20) | (0.29) | (0.33) |

Each row is a separate estimate of equation (4) with the stated outcome. Lower exposure is $\hat{\beta}+\hat{\rho}$, the response to a dry shock after the default in the bottom tercile of frozen non-bank reliance, with the state-specific coefficients averaged, higher exposure is $\hat{\beta}+\hat{\rho}+\hat{\theta}$, and the difference is $\hat{\theta}$ with its standard error. Log outcomes are multiplied by one hundred, so the entries are approximate percent changes. Wage income is log income from wages, defined for household-blocks with wage income. The three share rows sum to zero by construction and the essential row is reported as the implied value. Baseline is the mean share over the window. Standard errors clustered by district in parentheses. *, ** and *** denote significance at the 10, 5 and 1 percent levels.

of the good, from appliances and jewellery through clothing and footwear to the foods bought daily, the rank correlation with the estimated order is 0.63, and the rank correlation with the Engel elasticities of Appendix Table A4 is 0.55. The two orderings coincide at the top, where the categories are both durable and luxuries, and separate at education, milk and health, which are where the credit result lives. Appendix E discusses the rows that depart from the order, intoxicants among them, and the wet tail.

*Credit Supply and the Depth of Adjustment.* Panel A of Figure 3 is the centre of the paper. Where exposure was lower, the profile is shallow and it stops. Recreation falls 10 percent, restaurant meals 9, and health, education, milk and vegetables do not move. Where exposure was higher, the profile is deeper and it extends. Recreation falls 20 percent and restaurant

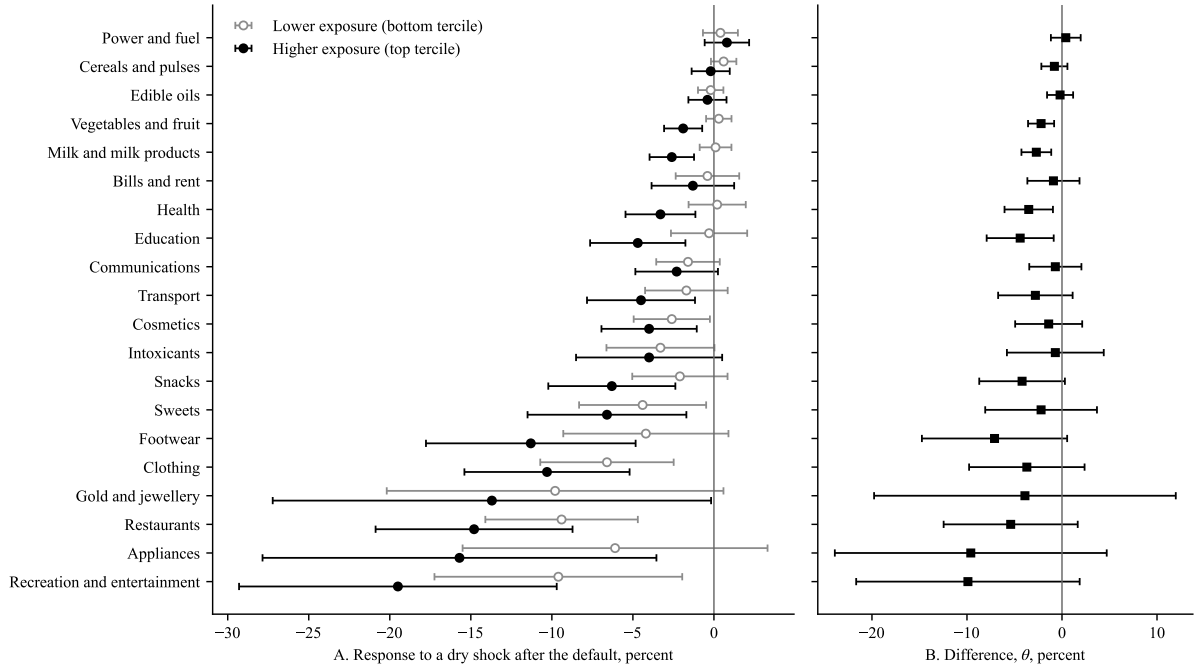



**Figure 3:** Category-Level Responses to a Dry Shock by Credit-Supply Exposure
Panel A plots, for each category, the response of rupees spent on the category to a dry shock after the default where exposure was lower, the bottom tercile of frozen non-bank reliance, and where it was higher, the top tercile, from separate estimates of equation (4) with the inverse hyperbolic sine of category rupees as the outcome, read as percent changes, in the three-column form of Section 3.2. Panel B plots the difference, $\hat{\theta}$. Categories are ordered by the all-household dry-tail response of Appendix Figure A4. Intervals are 95 percent, from standard errors clustered by district, over 2,196,418 household-blocks.

meals 15, and health spending falls 3.3 percent, education 4.7, milk 2.6 and vegetables 1.9. The protected index of the three named categories falls by 3.2 percent more where exposure was higher, with a standard error of 0.8 and a $q$-value across the rows of Tables 6 and 7 of 0.01. These are categories a household protects when it can, and this is prediction (4). Two features of the profile need separating, because the discretionary rows are roughly twice as deep where exposure was higher, and a larger income shock would also have produced that. The protected rows move from zero, which a shock of the same size would not have produced, and Table 5 finds the income shock the same size in both regimes within its precision. A shock close to twice as large, which that table cannot exclude, could in principle push a household past a buffer that the smaller shock left intact, so the distinction rests on the income test together with the placebo checks and the within-state identification of the main design, and not on the income test alone. The welfare claim of the paper rests on the second feature and not the first. Not every row is clean, since the difference between regimes is imprecise for the smallest categories, and gold and jewellery, bought in one household-month in thirty, has a difference of −3.9 percent with a standard error of 8.1, so its sign is not established, and because the survey's jewellery head is largely gold its cut may be a fall in saving rather than in consumption. Where exposure was higher, households exhaust their postponable spending and begin cutting what cannot be

postponed, and whether those cuts reverse is a question the panel can answer only within its own span, which we take up below.

*Borrowing and Other Margins.* Consumption is one margin of the portfolio, and Table 6 shows the others, with the same three columns and, beside each difference, its $q$-value.

**Table 6:** The Coping Portfolio: Borrowing, Assets and Labour

| | Baseline | Effect of a dry shock after the default | | | |
|---|---|---|---|---|---|
| | | Lower exposure | Higher exposure | Difference, $\theta$ | $q$ |
| *Panel A. Outstanding borrowing by source, share of households (pp), waves 13 to 18* | | | | | |
| From a bank | 18.4 | 0.71* | 0.37 | -0.34 | [0.54] |
| | | (0.42) | (0.56) | (0.69) | |
| From a non-bank finance company or dealer | 3.1 | 0.19 | -0.47* | -0.66** | [0.08] |
| | | (0.17) | (0.25) | (0.29) | |
| From a self-help group or microfinance lender | 5.6 | 0.21 | 0.43 | 0.22 | [0.54] |
| | | (0.27) | (0.37) | (0.43) | |
| From a moneylender | 8.6 | 0.33 | 0.97** | 0.64 | [0.30] |
| | | (0.31) | (0.43) | (0.44) | |
| From relatives or friends | 13.7 | 0.52 | 0.86 | 0.34 | [0.54] |
| | | (0.40) | (0.55) | (0.63) | |
| From shops (overdue accounts) | 21.4 | 0.44 | -0.07 | -0.51 | [0.54] |
| | | (0.49) | (0.67) | (0.76) | |
| *Panel B. Outstanding borrowing by purpose, share of households (pp), waves 13 to 18* | | | | | |
| For medical expenses | 1.9 | 0.11 | 0.32* | 0.21 | [0.45] |
| | | (0.12) | (0.18) | (0.20) | |
| To repay other debt | 4.6 | 0.08 | 0.41 | 0.33 | [0.45] |
| | | (0.20) | (0.29) | (0.32) | |
| *Panel C. Assets, gold and labour, share of households (pp)* | | | | | |
| Income from the sale of an asset | 2.6 | 0.14 | 0.26 | 0.12 | [0.63] |
| | | (0.22) | (0.32) | (0.41) | |
| Saved in gold in the last 120 days | 7.4 | 0.09 | -0.29 | -0.38 | [0.53] |
| | | (0.31) | (0.43) | (0.50) | |
| An additional member in work | | 0.22 | 0.37 | 0.15 | [0.54] |
| | | (0.17) | (0.29) | (0.31) | |
| Household-waves, Panels A and B | 880,588 (waves 13 to 18) | | | | |
| Household-waves, Panel C | 880,588, and 1,510,197 (waves 8 to 18, each wave against the one before) for the added member | | | | |

Each row is a separate estimate of equation (4) with the stated margin as $Y_{hdt}$, in the three-column form of Section 3.2, so the final coefficient column is $\hat{\theta}$. The outcome in every row is a binary indicator and the effect is a change in the share of households on that margin, not in how much they borrow or sell, since the survey records no loan amounts. Borrowing is a stock on the date of the interview and is measured over waves 13 to 18, with the unnamed-source box left out. Asset sales are income from the sale of an asset in the block, gold saving is the survey's indicator for having saved in gold over the preceding four months, and the added member is a rise in the household's count of members in work between waves. Baseline is the mean of the outcome over the stated window. $q$ is the sharpened false discovery rate $q$-value of Anderson (2008) across the twenty rows of this table and Table 7 together, with values below 0.01 shown as 0.01. Standard errors clustered by district in parentheses. *, ** and *** denote significance at the 10, 5 and 1 percent levels from the district-clustered errors.

One row survives the $q$-value at ten percent, non-bank borrowing, and the rest are imprecise. On the non-bank margin, which carries prediction (5), the share of households with a loan outstanding from a non-bank rises 0.19 percentage points after a dry shock where exposure was lower and falls 0.47 where it was higher, a difference of 0.66 points, a fifth of the margin's base, significant at five percent with a $q$-value of 0.08. This row is the most direct evidence in the survey that the treatment operates as intended, and it is also small, since the base is 3.1 percent of household-waves, so a fifth of it is one household in a hundred and fifty, and Section 5 returns to what that implies. Bank borrowing is imprecise, $-0.34$ with a standard error of 0.69,

and the interval runs from a fall two and a half times the non-bank fall to a rise, so the row neither confirms nor contradicts the supply-side evidence that banks held their lending. More households turn to a moneylender at the same time as the non-bank margin falls, since the share rises 0.97 percentage points where exposure was higher against 0.33 where it was lower, but the difference, 0.64 points with a standard error of 0.44, seven percent of the margin's base, is not distinguishable from zero and carries a $q$-value of 0.30, so the moneylender half of prediction (5) is suggestive at the precision the within-state design affords. The point estimates for borrowing from relatives and friends rise in both regimes, and the difference, 0.34 with a standard error of 0.63, is not distinguishable from zero, which is what Section 2 predicted for reciprocal insurance. Borrowing from shopkeepers, the largest single source of household credit in the survey, moves the wrong way, with a difference between regimes of $-0.51$ points and a standard error of 0.76 that leaves the sign open, and Appendix E sets out the two readings, trade credit cut off at the supplier or an overdue account that shrinks with purchases, that the data cannot separate. The remaining margins are estimated with the least precision. More households borrow for medical expenses, more borrow to repay a debt they already hold, more sell an asset, more send an additional member to work, and fewer put money into gold. All five move in the direction prediction (7) requires, none is individually distinguishable from zero, and none survives the $q$-value, so the escalation reading rests on five small estimates pointing the same way and we regard it as suggestive.

*Protected Spending and Debt Service.* Table 7 reports the protected categories and the one place the survey records rupees of debt. The protected index falls 3.2 percent more where exposure was higher, and each of its three components falls, with $q$-values of 0.01 to 0.07. The extensive margin moves with it, since the share of households spending anything on education in the block falls by 0.74 percentage points more where exposure was higher, on a base of 56 percent of household-blocks, and the share spending anything on health by 0.54 points on a base of 96 percent, each with a standard error of about two thirds its size, so the rupee cuts are consistent with households that stop buying and not only with smaller purchases at higher prices, though the extensive margins are imprecise on their own. Instalments paid on loans fall, imprecisely, by 4.9 percent after a dry shock where exposure was higher and rise, also imprecisely, by 1.8 percent where it was lower, a difference of 6.7 percent that is significant at ten percent, which is what a market that stopped making new loans in exposed districts should produce as the old ones run off, and the fall is in vehicle and consumer-durable instalments, the products non-banks made, and not in housing instalments.

*Transitions Between Borrowing States.* Formal incidence falling while informal incidence rises is suggestive, but by itself it establishes nothing about switching, since the households losing formal access and the households taking up informal credit could be different sets, and household fixed effects never link the two margins. Binary indicators can still answer the question if the outcome is the joint state. We classify every household-wave into one of four states, holding

**Table 7:** The Coping Portfolio: Protected Spending and Debt Service

| | Baseline | Effect of a dry shock after the default | | | |
|---|---|---|---|---|---|
| | | Lower exposure | Higher exposure | Difference, $\theta$ | $q$ |
| *Panel A. Protected categories, percent change in rupees, 2014 to 2019* | | | | | |
| Protected index (log of health, education and milk rupees) | | 0.0 | -3.2*** | -3.2*** | [0.01] |
| | | (0.5) | (0.7) | (0.8) | |
| Health | | 0.2 | -3.3*** | -3.5*** | [0.05] |
| | | (0.9) | (1.1) | (1.3) | |
| Education | | -0.3 | -4.7*** | -4.4** | [0.07] |
| | | (1.2) | (1.6) | (1.8) | |
| Milk and milk products | | 0.1 | -2.6*** | -2.7*** | [0.01] |
| | | (0.6) | (0.7) | (0.8) | |
| *Panel B. Debt service, percent change in rupees, 2014 to 2019* | | | | | |
| Instalments paid on loans (all EMIs) | | 1.8 | -4.9 | -6.7* | [0.24] |
| | | (2.1) | (3.1) | (3.9) | |
| Vehicle and durable instalments | | 2.4 | -6.9 | -9.3* | [0.24] |
| | | (3.0) | (4.4) | (5.5) | |
| Housing instalments | | 1.1 | -1.3 | -2.4 | [0.57] |
| | | (3.4) | (4.8) | (6.1) | |
| *Panel C. Extensive margin, share of households buying at all (pp), 2014 to 2019* | | | | | |
| Any spending on education in the block | 55.8 | -0.12 | -0.86* | -0.74 | [0.30] |
| | | (0.31) | (0.44) | (0.52) | |
| Any spending on health in the block | 95.9 | -0.08 | -0.62* | -0.54 | [0.30] |
| | | (0.22) | (0.32) | (0.38) | |
| Household-blocks | 2,196,418 | | | | |

Each row is a separate estimate of equation (4) with the stated outcome, in the three-column form of Section 3.2. The protected index is the log of rupees spent on health, education and milk together, named as a primary outcome before estimation, and the component rows and the instalment rows are the inverse hyperbolic sine of rupees on each, which keeps the household-blocks with no purchase, so every row runs over the full sample. Instalments are the survey's equated monthly instalment head and its vehicle, consumer-durable and housing parts. The extensive-margin rows are indicators for any spending on the category in the block. Log outcomes are reported as percent changes and indicators in percentage points of households. $q$ is the sharpened false discovery rate $q$-value across the twenty rows of Table 6 and this table together, with values below 0.01 shown as 0.01. Standard errors clustered by district in parentheses. *, ** and *** denote significance at the 10, 5 and 1 percent levels.

neither source, formal only, informal only, or both, and estimate the treatment effect on each of the twelve transitions between consecutive waves. Table 8 reports the result with the base rate of each transition beside it.

Some households do move, since among households holding formal credit alone the probability of holding informal credit alone a wave later rises by 2.7 percentage points on a base of 10 percent, with a standard error of 1.8, while the reverse transition, from informal credit alone to formal credit alone, is $-0.3$ points with a standard error of 0.7 on a base of 3.7 percent, an interval that admits a fall of nearly half and a rise of about a quarter, so the movement runs one way as far as the data can tell, which at this precision is not far. Households also leave formal credit for no credit at all, since the probability of going from formal credit alone to no loan from any source rises by 3.9 points on a base of 17 percent, significant at ten percent. Counting every exit from a state with a formal source, a dry shock where exposure was higher raises the share of formal borrowers who keep or find an informal source by 2.2 percentage points and the share who end with no loan by 2.7 points, and the difference, 0.5 with a standard error of 1.8, is not

**Table 8:** Transitions between Borrowing States

| State over borrowing sources | | | % of household-waves |
|---|---|---|---|
| *Panel A. Share of household-waves in each state, waves 13 to 18* | | | |
| Neither source | | | 51.3 |
| Formal only | | | 12.7 |
| Informal only | | | 24.3 |
| Both sources | | | 11.7 |
| *Panel B. Transitions between consecutive waves, conditional on the origin state: base rate and the difference, θ, in the effect of a dry shock after the default between higher and lower exposure* | | | |
| Transition | Base rate (%) | θ (pp) | Std. error |
| Neither source to formal only | 3.7 | -0.33 | (0.40) |
| Neither source to informal only | 11.3 | 0.60 | (0.67) |
| Neither source to both | 1.2 | 0.03 | (0.19) |
| Formal only to neither source | 16.5 | 3.90* | (2.25) |
| Formal only to informal only | 10.2 | 2.69 | (1.78) |
| Formal only to both | 8.7 | -0.77 | (1.41) |
| Informal only to neither source | 28.4 | 0.37 | (1.62) |
| Informal only to formal only | 3.7 | -0.30 | (0.69) |
| Informal only to both | 6.6 | -0.47 | (0.59) |
| Both to neither source | 7.7 | 1.32 | (1.22) |
| Both to formal only | 12.0 | -0.21 | (1.02) |
| Both to informal only | 19.7 | 1.67 | (1.32) |
| Loses formal access and keeps or finds an informal source | | 2.20** | (1.12) |
| Loses formal access and holds no loan from any source | | 2.66** | (1.31) |
| Difference between the two | | 0.47 | (1.76) |
| Household-wave pairs | | | 702,118 |

Panel A gives the unconditional distribution of household-waves across the four states over waves 13 to 18, with banks, non-bank finance companies, self-help groups, microfinance lenders and credit cards as formal sources and moneylenders, relatives and friends, shops, employers, chit funds and other sources as informal ones, and the unnamed-source box left out. Panel B reports, for each of the twelve transitions between consecutive waves of the same household, the share of the pairs that start in the origin state and make the transition, and the estimate of $\theta$ from equation (4) on those pairs with an indicator for the transition as the outcome, in percentage points of the pairs in the origin state. A pair is two consecutive interviews of the same household, and it takes its post-default indicator, its weather shock and its state-by-block effect from the later interview. The last three rows are the weighted sums of the $\theta$ estimates on the transitions out of the two states with a formal source into a state with an informal source and into the state with no source, in percentage points of the pairs that start with a formal source, and their difference. Standard errors clustered by district in parentheses. *, ** and *** denote significance at the 10, 5 and 1 percent levels.

established. Among the households that lose formal access after a shock, then, about as many stop borrowing as keep or find an informal source, and losing a formal lender is about as often rationing as redirection, which is prediction (6) in its defensible form. None of this establishes how much credit moved, since the survey records whether a household holds a loan and never its size, and a stock indicator cannot separate a household that was refused a new loan from one whose old loan ran off, which is why the paper speaks of losing access and never of being refused.

*Who Responds, and for How Long.* The design implies that $\theta$ should be concentrated where the income effect of the shock is, and Appendix Table A5 estimates it by exposure group for the four main margins. The discretionary and protected-index cuts are largest for agricultural labourers and small farmers, present for construction labourers, smaller and imprecise for the remaining groups, and null for salaried and retired households on every margin, which is the pattern the placebo logic requires and the all-household estimates cannot show. The fall in non-bank borrowing is spread across every non-placebo group, including the two with no income effect, one of which, the other self-employed, is where the non-bank lenders' customers were, a reminder that the borrowing row records the lender's withdrawal and not the household's need. Appendix Table A6 splits the estimates by rural and urban residence. The point estimates are

larger in the countryside on every spending margin and on the moneylender margin, and no rural-urban difference is distinguishable from zero. Appendix Table A8 adds the shock in the previous block and the block before that to equation (4). The point estimates for the health and education cuts and for the fall in non-bank borrowing one block after a dry shock are about three fifths of the contemporaneous ones and no longer distinguishable from zero, those two blocks after are close to zero, and the milk cut is gone within a block. Within the span of these data, then, the protected cuts and the loss of the lender fade within eight months as far as the estimates can tell, and whether a child kept out of school for a term returns is beyond what the survey records.

### 4.3 Robustness

Appendix Tables A9 to A12 collect the checks, and Appendix B discusses the exposure, index and inference checks in detail. The exposure measure can be entered as the standardised share, built from rupees, built on the 2015 to 2016 window, built with gold loans, built with fintech lenders counted as non-banks, or built from the vehicle, property and housing lending of non-banks alone, and every alternative moves the estimates by less than a standard error, while the share standardised to unit variance gives about half the tercile contrast, as it should when the terciles sit 2.3 standard deviations apart. Adding the district's urban share interacted with the shock and the post-default indicator shrinks $\theta$ by a sixth or less, comparing the top tercile with the middle roughly halves it, and comparing the middle with the bottom gives about half again, as it should if the middle tercile also lost non-bank lending. The design rows answer the considerations of Section 2.3. The pooled design, in which the weather response and its change at the default are common across states so that the cross-state variation enters too, gives $-0.64$ with a standard error of 0.26, a quarter larger than the within-state estimate and significant at five percent, and interacting every district characteristic of Appendix Table A3 with the shock and the post-default indicator within the main design gives $-0.40$ with a standard error of 0.36, so the within-state headline is the conservative reading and the covariates row is its lower edge. Excluding the four states that waived farm loans in 2018 leaves the estimate at $-0.56$, excluding farming households at $-0.47$, and applying the survey weights at $-0.54$. Thresholds over the full record or within calendar month, winsorising the index, the monsoon months alone, the cell-level intensity share, the count of dry months, the blocks with a month in each tail dropped, the straddling blocks coded by majority, the pre-period block with the largest coefficient dropped, the de Chaisemartin and D'Haultfœuille (2020) estimator, the districts present since 2014, and the sample without the demonetisation blocks all return estimates within a standard error of the main one, while the three- and six-month accumulation windows give larger weather coefficients and less precise credit interactions. Two rows do not fit the rest, as the wet tail produces a weather-only response of the same size as the dry tail but no credit interaction, $-0.17$ with a standard error of 0.32, so the contraction changes how households absorb a dry shock and not

how they absorb a wet month, which the poorer measurement of the wet tail and the national clustering of the wet months of 2019 may explain and which we cannot resolve. And at monthly frequency the dry-shock coefficient on the discretionary share is between a quarter and a third of the block estimate, and the credit interaction a fifth, which is close to what a monthly indicator should give as a proxy for a block that holds one and a half dry months in four on average, and is discussed in Appendix B.

The falsification tests of Appendix Table A12 return nothing in twelve of fourteen rows. The salaried placebo is null on income in both tails, on the budget share, and on the triple difference for non-bank borrowing, and the retired placebo is null on the triple difference for the discretionary share. Future weather does not predict current spending, and district originations of loans of fifty thousand rupees or more, which are not a consumption-smoothing instrument, do not respond to the dry tail in the bureau. Dating the default a year early and ending the sample in August 2018 gives a $\theta$ of 0.07 with a standard error of 0.37, the household counterpart of the placebo date in Table 1. The probability of a household leaving the cleaned panel described in Appendix C in the wave after a dry shock, where leaving means not being interviewed or returning an incomplete or flagged block, falls by 0.6 percentage points on a base of 12.7 percent, significant at ten percent and with the wrong sign for attrition to inflate the estimates, and the triple difference on the same exit indicator is $-0.18$ with a standard error of 0.57. The exceptions are the retired placebo in the wet tail and that attrition row. Appendix Table A11 reports the main $\theta$ under spatial standard errors at two and five hundred kilometres, clustering by homogeneous region and by state with the wild cluster bootstrap, and the exposure permutation within state and across all districts. It is significant at ten percent under district clustering and under the permutation across all districts, and its $p$-value lies between ten and fifteen percent under the spatial, regional and state procedures, so the headline is a ten-percent result at the precision the within-state design affords, while the pooled design that uses the cross-state variation too is significant at five percent under district clustering.

# 5 Interpretation

## 5.1 The Household's Own Lender and the District's Credit Market

The direct margin is small, and the paper's reading has to be consistent with that. Three percent of household-waves hold a non-bank loan, the contraction removes a fifth of them after a dry shock, and the discretionary share falls by 0.52 points across all households in the district. If the spending response were confined to the households that lost their own non-bank loan, each of them would have to cut its discretionary share by tens of points, which no household does, so most of the response is in households that did not hold a non-bank loan. The estimand of equation (4) is a district-level one, and this is what it looks like. Breza and Kinnan (2021) found the same structure when microfinance contracted, since consumption fell for households that

had never borrowed from a microfinance lender, and traced it to local wages and demand. Here the candidates are the merchants and employers whose working capital came from the same lenders, and the shop credit that moved the wrong way in Table 6 is consistent with that, and the small-ticket lending that answered dry months before the default and stopped answering them after (Table 2), which reaches a household without appearing in a survey that asks about outstanding loans once a wave. The microfinance lenders that share the non-banks' funding are a third candidate in principle, and Table 6 gives no sign that borrowing from them fell, so we do not lean on it. Table 5 gives no sign of the reading in which the district channel is a larger income shock, since income fell by the same amount in both regimes within its precision, which cannot exclude a shock close to twice as large, so what the district lost was slack, the credit and advances that let a household wait, and not earnings, and the data cannot say more than that about which local channel carried the rest. The withdrawal of formal credit changed how the households of a district absorbed a dry shock, and the household's own lender is one of the ways it did.

A second question the tables raise is why the households the non-banks had been serving did not move to banks, which held their lending. The transitions in Table 8 speak against one reading, since movement from informal credit back to formal credit does not rise, with an interval that admits a fall of nearly half and a rise of about a quarter on its base, so displaced households did not reach a bank by way of a moneylender within the panel as far as the data can tell. Whether non-bank borrowers are not bankable at all, or whether the adjustment takes longer than the four blocks the panel allows, the data cannot distinguish, and Appendix E sets out the two readings and the one clue the bureau offers.

### 5.2 Credit, Transfers and the Literature

Government transfers reached 1.1 percent of household-months in 2014 and 60.6 percent by 2019, over exactly the years in which formal credit expanded, so a reader will ask how the two buffers compare. Appendix E reports a comparison on the exposed occupations with the district's formal originations per person and its transfer coverage entered together on one scale, each frozen before the panel. The comparison is descriptive and sits outside the design of equation (4), because no transfer rollout can be frozen before the shock the way the non-bank share can, and what it supports is a ranking that is suggestive and not established. The discretionary-share cut is 0.31 percentage points smaller per standard deviation of formal credit in the district and 0.12 smaller per standard deviation of transfer coverage, and the transfer estimate is distinguishable neither from zero nor from the credit estimate. Inside the design, the credit effect is smaller for households that received a transfer in the preceding wave, $-0.66$ against $-0.97$, and the difference is not established either.

Against the literature, the paper's magnitudes sit where its design puts them. The access experiments find that credit blunts the effect of a shock on consumption or on the expenses

households forgo (Jack and Suri, 2014; Lane, 2024; Suri et al., 2021), and the extensive-margin rows of Table 7, households that stop buying health care or schooling at all, are the closest counterpart to the expenses forgone that Suri et al. (2021) measure. Breza and Kinnan (2021) find that the Andhra Pradesh contraction, which removed roughly half of microcredit in exposed districts, lowered the casual wage by about 2 percent and monthly household consumption by under 3 percent, through wages and demand. Ours removed an eighth of the non-bank originations in exposed districts, a smaller cut to a far larger market, and moved the composition of the basket. We do not report a consumption multiplier of their kind, because the survey records no loan amounts and the spending response is a reallocation, so the numerator and the denominator would be in different units. Cramer et al. (2025) find on the same bureau that the response of non-bank lending to a weather shock falls after the default, and Table 2 finds the same on district exposure rather than lender type, so the household results above are what that fall looks like from the other side of the transaction, and Appendix Table A13 collects the published numbers.

### 5.3 Limitations

The limits of the design deserve one place of their own. The survey records whether a household holds a loan, never how much, and as a stock rather than a flow, so every borrowing result is a count of households and cannot separate a refused loan from one that ran off. The survey's non-bank category covers a finance company, a fintech lender and a dealer selling on instalment alike, so the lender-specific claims rest on the bureau. The bureau records no household, so the direct channel from a district's lost lending to its households is inferred and not observed, and the district channels that must carry most of the effect are named but not separated. Income and spending are reconciled on screen and income is smoothed in reporting, so the income effect is attenuated and the spending response is a reconciled number. The survey records rupees and not quantities, so the essential rows after a dry shock cannot separate protected consumption from higher prices, and the extensive margin is the one place where the distinction can be seen. The control group is the surveyed three fifths of the bottom tercile, the two groups were not on parallel paths in lending before the default, and the dry months after the default fell more often on control districts than on exposed ones, which thins the identifying cell. The treatment is a dry month inside a four-month block and not a failed season, and the headline is identified within the ten states that hold both terciles, where it is significant at ten percent, so the paper's precision is that of a within-state design and the pooled estimate that is significant at five percent rests on cross-state variation as well. The post-default window is sixteen months, so persistence can be followed for two blocks and no longer. The panel under-represents the poor, for whom the mechanism should be strongest, so the estimates apply to the surveyed population and are likely attenuated. And the wet tail is poorly measured and arrived in national clusters after the default, so the absence of a credit interaction in that tail is a null we cannot interpret. The

paper's claims are the ones that survive these limits, that the same dry shock was absorbed differently where formal credit had contracted more, on margins the survey measures well, and that the difference ran through the district's credit market and not only through the household's own lender.

# 6 Conclusion

We use the default of IL&FS as a natural experiment to measure what a contraction in formal credit supply does to the way the households of a district absorb a dry shock. A failure in infrastructure finance closed the wholesale funding market for India's non-bank lenders within weeks, and the districts whose borrowers had relied on those lenders lost an eighth of their non-bank lending for reasons that had nothing to do with their rainfall. Setting the formal retail lending that reaches the bureau beside a panel that observes consumption by category, borrowing by source and purpose, asset sales, gold saving and labour supply for the same households, we can see both sides of the transaction that the credit-supply literature has so far observed from one side only.

Households facing a weather shock do not first borrow, and they do not first cut back, because what they do first is rearrange. Total spending is close to unchanged while the composition of the basket moves substantially, and what goes first runs from the most postponable to the least, which is why the consumption-smoothing literature has so often found nothing when it regressed spending levels on rainfall. Whether a household stops at rearranging depends on the credit market around it. The same dry shock took the same share of income in districts more and less exposed to the contraction, and where exposure was higher it more than doubled the depth of the discretionary cut and extended it into spending that households otherwise protect, health, education and a staple such as milk, with the extension and not the doubling as the evidence that credit, and not a larger shock, is at work. On the borrowing margins, fewer households hold a non-bank loan, more hold one from a moneylender, though imprecisely, instalments paid on vehicle and durable loans fall, and among the households that lose formal access about as many stop borrowing as keep or find an informal source. The direct margin is small, because few households borrow from a non-bank, so most of the spending response runs through the district's credit market, as it did when microcredit contracted a decade earlier, and the data name the candidate channels without separating them.

The policy reading is narrower than it might appear. Credit does not prevent a dry shock from being felt, since the income loss was the same in both regimes. What the credit market around a household determines is how far along the coping portfolio it must travel before the gap closes, and within the span of these data the cuts to health and education and the loss of the lender had faded within two blocks of a dry shock, with what happens after that beyond the panel. The episode also shows that a funding shock to non-bank lenders, an event studied so

far on the balance sheets of banks and funds, has a household incidence, and that the incidence falls on the margins of the budget that regulators of that funding would least want it to.

Two additions would most improve on this design, the first being loan amounts, which would turn every extensive-margin result into an intensive one and would let the direct and the district channels be separated. And a credible instrument for non-bank reliance, in place of a frozen baseline share, would let the design survive the objection that districts relying on non-bank lenders before the default were already different. The estimates as they stand apply to the surveyed population, which under-represents the households for whom the mechanism should be strongest, and the frozen share measures who was lending in the year before the default, so what we estimate is the short-run consequence of losing a lender that was already present, not the long-run consequence of never having had one (Burgess and Pande, 2005), and the two need not share a sign.

# References


Abraham, R. and A. Shrivastava (2022). "How Comparable are India's Labour Market Surveys?" *Centre for Sustainable Employment Working Paper*, Azim Premji University.

Acharya, V. V., H. Khandwala and T. S. Öncü (2013). "The Growth of a Shadow Banking System in Emerging Markets: Evidence from India." *Journal of International Money and Finance* 39, 207–230.

Anderson, M. L. (2008). "Multiple Inference and Gender Differences in the Effects of Early Intervention: A Reevaluation of the Abecedarian, Perry Preschool, and Early Training Projects." *Journal of the American Statistical Association* 103(484), 1481–1495.

Banerjee, A. V. and E. Duflo (2007). "The Economic Lives of the Poor." *Journal of Economic Perspectives* 21(1), 141–167.

Bhardwaj, A. and A. Javadekar (2024). "How Does Bank Lending to Non-Banks Affect Credit Allocation and Systemic Risk?" Working paper.

Breza, E. and C. Kinnan (2021). "Measuring the Equilibrium Impacts of Credit: Evidence from the Indian Microfinance Crisis." *Quarterly Journal of Economics* 136(3), 1447–1497.

Browning, M. and T. F. Crossley (2009). "Shocks, Stocks, and Socks: Smoothing Consumption over a Temporary Income Loss." *Journal of the European Economic Association* 7(6), 1169–1192.

Buchak, G., G. Matvos, T. Piskorski and A. Seru (2018). "Fintech, Regulatory Arbitrage, and the Rise of Shadow Banks." *Journal of Financial Economics* 130(3), 453–483.

Burgess, R. and R. Pande (2005). "Do Rural Banks Matter? Evidence from the Indian Social Banking Experiment." *American Economic Review* 95(3), 780–795.

Callaway, B., A. Goodman-Bacon and P. H. C. Sant'Anna (2024). "Difference-in-Differences with a Continuous Treatment." NBER Working Paper 32117.

Cameron, A. C., J. B. Gelbach and D. L. Miller (2008). "Bootstrap-Based Improvements for Inference with Clustered Errors." *Review of Economics and Statistics* 90(3), 414–427.

Carter, M. R. and C. B. Barrett (2006). "The Economics of Poverty Traps and Persistent Poverty: An Asset-Based Approach." *Journal of Development Studies* 42(2), 178–199.

Chen, J. and J. Roth (2024). "Logs with Zeros? Some Problems and Solutions." *Quarterly Journal of Economics* 139(2), 891–936.

Chodorow-Reich, G. (2014). "The Employment Effects of Credit Market Disruptions: Firm-Level Evidence from the 2008–9 Financial Crisis." *Quarterly Journal of Economics* 129(1), 1–59.

Conley, T. G. (1999). "GMM Estimation with Cross Sectional Dependence." *Journal of Econometrics* 92(1), 1–45.

Corno, L., N. Hildebrandt and A. Voena (2020). "Age of Marriage, Weather Shocks, and the Direction of Marriage Payments." *Econometrica* 88(3), 879–915.

Cramer, K. F., P. Ghosh, N. Kulkarni and N. Vats (2025). "Shadow Banks on the Rise: Evidence Across Market Segments." Financial Markets Group Discussion Paper 934, London School of Economics.

de Chaisemartin, C. and X. D'Haultfœuille (2020). "Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects." *American Economic Review* 110(9), 2964–2996.

Deaton, A. (1991). "Saving and Liquidity Constraints." *Econometrica* 59(5), 1221–1248.

Deaton, A. (1997). *The Analysis of Household Surveys: A Microeconometric Approach to Development Policy*. Baltimore: Johns Hopkins University Press.

Dell, M., B. F. Jones and B. A. Olken (2014). "What Do We Learn from the Weather? The New Climate-Economy Literature." *Journal of Economic Literature* 52(3), 740–798.

Demont, T. (2020). "Coping with Shocks: The Impact of Self-Help Groups on Migration and Food Security." AMSE Working Paper 2020-16, Aix-Marseille School of Economics.

Drèze, J. and A. Somanchi (2024). "Weighty Evidence? Poverty Estimation with Missing Data." Working paper.

Dubey, T. S. and A. Purnanandam (2023). "Can Cashless Payments Spur Economic Growth?" Working paper.

Fafchamps, M., C. Udry and K. Czukas (1998). "Drought and Saving in West Africa: Are Livestock a Buffer Stock?" *Journal of Development Economics* 55(2), 273–305.

Ghosh, P. and N. Vats (2024). "Safety Nets, Credit, and Investment: Evidence from a Guaranteed Income Program." Working paper, SSRN 4265112.

Goldsmith-Pinkham, P., I. Sorkin and H. Swift (2020). "Bartik Instruments: What, When, Why, and How." *American Economic Review* 110(8), 2586–2624.

Gopal, M. and P. Schnabl (2022). "The Rise of Finance Companies and FinTech Lenders in Small Business Lending." *Review of Financial Studies* 35(11), 4859–4901.

Government of India, Ministry of Finance (2020). "Financial Fragility in the NBFC Sector." Chapter 8 of *Economic Survey 2019–20*, Volume I. New Delhi.

Huber, K. (2018). "Disentangling the Effects of a Banking Crisis: Evidence from German Firms and Counties." *American Economic Review* 108(3), 868–898.

Imbert, C. and J. Papp (2015). "Labor Market Effects of Social Programs: Evidence from India's Employment Guarantee." *American Economic Journal: Applied Economics* 7(2), 233–263.

Irani, R. M., R. Iyer, R. R. Meisenzahl and J.-L. Peydró (2021). "The Rise of Shadow Banking: Evidence from Capital Regulation." *Review of Financial Studies* 34(5), 2181–2235.

Jack, W. and T. Suri (2014). "Risk Sharing and Transactions Costs: Evidence from Kenya's Mobile Money Revolution." *American Economic Review* 104(1), 183–223.

Jayachandran, S. (2006). "Selling Labor Low: Wage Responses to Productivity Shocks in Developing Countries." *Journal of Political Economy* 114(3), 538–575.

Jensen, R. (2000). "Agricultural Volatility and Investments in Children." *American Economic Review* 90(2), 399–404.

Kaur, S. (2019). "Nominal Wage Rigidity in Village Labor Markets." *American Economic Review* 109(10), 3585–3616.

Kazianga, H. and C. Udry (2006). "Consumption Smoothing? Livestock, Insurance and Drought in Rural Burkina Faso." *Journal of Development Economics* 79(2), 413–446.

Khwaja, A. I. and A. Mian (2008). "Tracing the Impact of Bank Liquidity Shocks: Evidence from an Emerging Market." *American Economic Review* 98(4), 1413–1442.

Kulkarni, N., K. M. Neelima and S. Sinha (2023). "Anatomy of a Liquidity Shock on Non-banks." Working paper, August 2023.

Lane, G. (2024). "Adapting to Climate Risk with Guaranteed Credit: Evidence from Bangladesh." *Econometrica* 92(2), 355–386.

Maccini, S. and D. Yang (2009). "Under the Weather: Health, Schooling, and Economic Consequences of Early-Life Rainfall." *American Economic Review* 99(3), 1006–1026.

Malaeb, B., J. Golan and K. Imai (2026). "Consumption Smoothing and Risk Coping Mechanisms: Evidence from Rural India." *Review of Income and Wealth*.

Morduch, J. (1995). "Income Smoothing and Consumption Smoothing." *Journal of Economic Perspectives* 9(3), 103–114.

Muralidharan, K., P. Niehaus and S. Sukhtankar (2016). "Building State Capacity: Evidence from Biometric Smartcards in India." *American Economic Review* 106(10), 2895–2929.

Paxson, C. H. (1992). "Using Weather Variability to Estimate the Response of Savings to Transitory Income in Thailand." *American Economic Review* 82(1), 15–33.

Rambachan, A. and J. Roth (2023). "A More Credible Approach to Parallel Trends." *Review of Economic Studies* 90(5), 2555–2591.

Rosenzweig, M. R. and K. I. Wolpin (1993). "Credit Market Constraints, Consumption Smoothing, and the Accumulation of Durable Production Assets in Low-Income Countries: Investments in Bullocks in India." *Journal of Political Economy* 101(2), 223–244.

Sengupta, R., L. L. Son and H. Vardhan (2022). "A Study of the Non-Banking Finance Companies in India." Working Paper WP-2022-009, Indira Gandhi Institute of Development Research, Mumbai.

Shah, M. and B. M. Steinberg (2017). "Drought of Opportunities: Contemporaneous and Long-Term Impacts of Rainfall Shocks on Human Capital." *Journal of Political Economy* 125(2), 527–561.

Sinha Roy, S. and R. van der Weide (2025). "Estimating Poverty for India after 2011 Using Private-Sector Survey Data." *Journal of Development Economics* 172, 103386.

Somanchi, A. (2021). "Missing the Poor, Big Time: A Critical Assessment of the Consumer Pyramids Household Survey." SocArXiv working paper.

Suri, T., P. Bharadwaj and W. Jack (2021). "Fintech and Household Resilience to Shocks: Evidence from Digital Loans in Kenya." *Journal of Development Economics* 153, 102697.

Townsend, R. M. (1994). "Risk and Insurance in Village India." *Econometrica* 62(3), 539–591.

Udry, C. (1994). "Risk and Insurance in a Rural Credit Market: An Empirical Investigation in Northern Nigeria." *Review of Economic Studies* 61(3), 495–526.

Vicente-Serrano, S. M., S. Beguería and J. I. López-Moreno (2010). "A Multiscalar Drought Index Sensitive to Global Warming: The Standardized Precipitation Evapotranspiration Index." *Journal of Climate* 23(7), 1696–1718.

# Appendix

## A Additional Tables and Figures

This appendix collects the descriptive tables and figures referred to in Sections 2 and 3, the by-group, residence, coefficient and persistence tables referred to in Section 4, and the robustness, inference, falsification and benchmark tables referred to in Sections 4 and 5.

**Table A1:** The Sample, by Exposure Group

| | Sample | | Monthly | Share of household-waves, %, waves 13 to 18 | | | Budget shares, % | | |
|---|---|---|---|---|---|---|---|---|---|
| Exposure group | Households | Rural, % | income, Rs | Asset sale | Non-bank loan | SHG or MFI loan | Food | Discret. | Essential |
| Agricultural labourer | 7,332 | 93.0 | 11,956 | 0.7 | 2.1 | 10.6 | 52.0 | 26.9 | 49.3 |
| Small farmer | 17,115 | 90.9 | 14,329 | 1.9 | 2.6 | 8.8 | 51.1 | 25.7 | 50.8 |
| Organised farmer | 8,127 | 83.7 | 20,787 | 3.9 | 3.4 | 6.2 | 48.4 | 25.4 | 51.8 |
| Self-employed, outdoor | 4,362 | 22.7 | 20,614 | 3.3 | 3.9 | 4.9 | 47.2 | 27.5 | 50.0 |
| Construction labourer | 28,912 | 42.6 | 12,515 | 2.8 | 2.8 | 8.1 | 53.5 | 26.4 | 50.8 |
| Casual labourer, indoor | 17,415 | 22.1 | 13,019 | 1.5 | 3.0 | 7.0 | 51.5 | 26.7 | 50.8 |
| Self-employed, other | 41,706 | 14.3 | 21,068 | 3.1 | 3.6 | 5.1 | 47.1 | 27.3 | 50.6 |
| Other (head not in the labour force) | 12,799 | 26.1 | 18,229 | 2.7 | 2.4 | 4.8 | 49.9 | 26.0 | 51.3 |
| *Salaried* | 46,761 | 15.0 | 24,714 | 2.7 | 3.5 | 3.7 | 46.0 | 27.8 | 50.1 |
| *Retired* | 26,031 | 20.0 | 24,566 | 3.1 | 2.2 | 3.1 | 46.6 | 26.1 | 51.6 |

Analysis panel, January 2014 to December 2019. Groups are built from the occupation of the household head at first observation, crossed with whether the head's industry is outdoor work, as Appendix C describes, and held fixed thereafter. Income is the mean of reported monthly household income across all sources. The three household-wave columns are shares of household-waves over waves 13 to 18 reporting income from the sale of an asset in the block, a loan outstanding from a non-bank finance company or dealer, and a loan outstanding from a self-help group or microfinance lender. Budget shares are means across household-blocks. The two italicised groups are placebos, in that their income does not depend on the monsoon.

**Table A2:** District-Blocks and Districts in Each Cell of the Design

| | Before the default | | | | After the default | | | |
|---|---|---|---|---|---|---|---|---|
| | No dry month | | A dry month | | No dry month | | A dry month | |
| Exposure | Blocks | Districts | Blocks | Districts | Blocks | Districts | Blocks | Districts |
| Top tercile | 1,360 | 206 | 1,524 | 204 | 495 | 204 | 329 | 187 |
| Middle tercile | 1,286 | 190 | 1,374 | 189 | 485 | 186 | 275 | 150 |
| Bottom tercile | 816 | 122 | 892 | 122 | 251 | 116 | 237 | 101 |

Counts over the 518 analysis districts and the eighteen four-month calendar blocks of 2014 to 2019, with a block counted as having a dry month if the dry tail fires in any of its months. Blocks are the calendar blocks of reference months, so the count of household-blocks in any cell is the sum over its households. The identifying cell for $\theta$ is the top tercile with a dry month after the default.

**Table A3:** District Characteristics by Pre-Default Non-Bank Reliance

| | Bottom tercile | Middle tercile | Top tercile | Top minus bottom | s.e. |
|---|---|---|---|---|---|
| *Panel A. Districts in the household panel with a weather index (the analysis sample)* | | | | | |
| Non-bank share of loans | 0.199 | 0.344 | 0.511 | 0.312*** | (0.007) |
| Non-bank share of rupees | 0.191 | 0.272 | 0.311 | 0.119*** | (0.006) |
| Non-bank rupees per person | 3,383 | 3,542 | 5,839 | 2,456*** | (588) |
| Loans per 1,000 people | 127.0 | 47.0 | 68.3 | -58.7*** | (11.2) |
| Rupees per person | 18,810 | 13,301 | 19,130 | 320 | (2,137) |
| Rupees per loan | 190,776 | 275,666 | 259,053 | 68,277*** | (8,631) |
| Urban share, 2011 (%) | 25.7 | 20.9 | 34.0 | 8.3*** | (2.3) |
| Population, 2011 (millions) | 1.94 | 2.03 | 2.55 | 0.60*** | (0.17) |
| Dry-tail months, 2014 to 2019 (%) | 19.8 | 18.0 | 18.7 | -1.1* | (0.5) |
| Wet-tail months, 2014 to 2019 (%) | 19.5 | 20.2 | 21.0 | 1.4*** | (0.4) |
| Districts | 122 | 190 | 206 | | |
| *Panel B. All bureau districts* | | | | | |
| Non-bank share of loans | 0.163 | 0.344 | 0.510 | 0.347*** | (0.007) |
| Non-bank share of rupees | 0.148 | 0.275 | 0.312 | 0.164*** | (0.007) |
| Non-bank rupees per person | 2,350 | 3,341 | 5,669 | 3,319*** | (538) |
| Loans per 1,000 people | 88.8 | 44.3 | 66.2 | -22.6*** | (8.6) |
| Rupees per person | 14,652 | 12,537 | 18,547 | 3,895** | (1,846) |
| Rupees per loan | 266,997 | 277,461 | 259,440 | -7,557 | (10,425) |
| Urban share, 2011 (%) | 25.6 | 20.3 | 33.3 | 7.7*** | (2.1) |
| Population, 2011 (millions) | 1.26 | 1.91 | 2.49 | 1.23*** | (0.15) |
| Dry-tail months, 2014 to 2019 (%) | 20.4 | 18.1 | 18.7 | -1.8*** | (0.5) |
| Wet-tail months, 2014 to 2019 (%) | 20.0 | 20.3 | 21.0 | 1.0*** | (0.3) |
| In the household panel (%) | 60.6 | 89.2 | 96.3 | 35.7*** | (3.6) |
| Districts | 213 | 213 | 214 | | |

Means by tercile of the non-bank and housing-finance share of loans originated from June 2017 to May 2018. Panel A covers the 518 districts with an exposure share, a weather index and surveyed households, over which every household estimate runs, and Panel B all 640 bureau districts. Loans and rupees are over the same twelve months, so they are annual rates. Population and the urban share are from the 2011 Census. Dry-tail and wet-tail months are the share of months in 2014 to 2019 in each tail of the district's own water-balance history. The standard error is for the difference between the top and bottom terciles, from a Welch test. *, ** and *** denote significance at the 10, 5 and 1 percent levels.

**Table A4:** The Consumption Basket

| Block | Category | Budget share, % | Households buying, % | Engel elasticity |
|---|---|---|---|---|
| Essential | Power and fuel | 15.21 | 98.7 | 1.25 |
| | Cereals and pulses | 13.35 | 99.9 | 0.46 |
| | Milk and milk products | 7.94 | 97.7 | 1.04 |
| | Vegetables and fruit | 7.48 | 99.9 | 0.78 |
| | Edible oils | 3.34 | 98.4 | 0.49 |
| | Health | 2.45 | 91.1 | 1.06 |
| | Bills and rent | 0.94 | 34.2 | 1.67 |
| | *Block total* | *50.71* | | |
| Discretionary | Toiletries and cosmetics | 6.02 | 100.0 | 0.88 |
| | Communications | 4.35 | 98.1 | 1.10 |
| | Intoxicants | 3.14 | 70.2 | 0.70 |
| | Education | 3.13 | 47.1 | 1.78 |
| | Clothing | 2.96 | 32.5 | 1.23 |
| | Transport | 2.54 | 83.4 | 0.93 |
| | Restaurants | 1.75 | 65.2 | 1.35 |
| | Footwear | 0.83 | 27.9 | 1.24 |
| | Snacks | 0.69 | 80.2 | 0.97 |
| | Sweets | 0.57 | 47.0 | 1.06 |
| | Recreation and entertainment | 0.42 | 18.3 | 1.95 |
| | Gold and jewellery | 0.18 | 3.1 | 2.49 |
| | Appliances | 0.15 | 3.6 | 2.10 |
| | *Block total* | *26.72* | | |
| Residual | Other food, instalments, miscellaneous | 22.57 | | |

Analysis panel, January 2014 to December 2019. Budget share is the mean share of total monthly expenditure. The third column is the share of household-months in which the category is bought at all, and the discretionary categories are statistically demanding because of it, since appliances and gold appear in fewer than one household-month in twenty. The recreation category is the survey's recreation head, which contains entertainment. The Engel elasticity is from a regression of the category's log spending on log total expenditure across households, with an elasticity above one marking a category whose share rises with the budget. The residual row is the part of total expenditure outside the twenty categories.

**Table A5:** The Credit Contraction by Exposure Group

| Exposure group, $\theta$ | Discretionary share (pp) | Protected index (%) | Non-bank borrowing (pp) | Moneylender borrowing (pp) |
|---|---|---|---|---|
| Agricultural labourer | -1.30* | -7.8** | -0.57 | 1.19 |
| | (0.69) | (3.1) | (0.36) | (1.00) |
| Small farmer | -1.06** | -6.3*** | -0.73** | 0.91 |
| | (0.52) | (2.4) | (0.29) | (0.79) |
| Organised farmer | -0.65 | -4.1 | -0.97** | 0.47 |
| | (0.66) | (2.9) | (0.43) | (0.95) |
| Self-employed, outdoor | -0.73 | -4.5 | -1.14** | 0.76 |
| | (0.88) | (3.9) | (0.58) | (1.34) |
| Construction labourer | -0.81* | -5.2*** | -0.73*** | 0.93 |
| | (0.43) | (1.9) | (0.26) | (0.62) |
| Casual labourer, indoor | -0.65 | -3.7 | -0.81** | 0.69 |
| | (0.49) | (2.3) | (0.32) | (0.72) |
| Self-employed, other | -0.45 | -3.2* | -0.89*** | 0.85 |
| | (0.37) | (1.7) | (0.25) | (0.55) |
| Other (head not in the labour force) | -0.33 | -2.4 | -0.57 | 0.49 |
| | (0.72) | (3.2) | (0.45) | (1.06) |
| *Salaried* | -0.08 | -0.5 | 0.07 | 0.07 |
| | (0.37) | (1.7) | (0.37) | (0.50) |
| *Retired* | -0.07 | -0.2 | 0.03 | -0.10 |
| | (0.39) | (1.8) | (0.42) | (0.52) |

Each cell is $\theta$ from equation (4) estimated on the stated exposure group alone, with the units and windows of Tables 6 and 7. Italicised groups are placebos. Standard errors clustered by district in parentheses. *, ** and *** denote significance at the 10, 5 and 1 percent levels.

**Table A6:** The Credit Contraction by Rural and Urban Residence

| Margin, $\theta$ by residence | Rural households | Urban households | Rural minus urban |
|---|---|---|---|
| Discretionary share (pp) | -0.58 | -0.47 | -0.11 |
| | (0.39) | (0.43) | (0.58) |
| Protected index (percent) | -3.6 | -2.2 | -1.4 |
| | (2.6) | (3.0) | (3.9) |
| Health spending (percent) | -3.7 | -2.4 | -1.4 |
| | (2.9) | (3.3) | (4.4) |
| Education spending (percent) | -5.0 | -3.6 | -1.4 |
| | (3.5) | (3.7) | (5.0) |
| Non-bank borrowing (pp) | -0.75* | -0.54 | -0.20 |
| | (0.45) | (0.39) | (0.60) |
| Moneylender borrowing (pp) | 0.98 | 0.36 | 0.63 |
| | (0.62) | (0.57) | (0.85) |
| Income from the sale of an asset (pp) | 0.07 | 0.25 | -0.18 |
| | (0.49) | (0.35) | (0.60) |
| An additional member in work (pp) | 0.22 | 0.10 | 0.12 |
| | (0.39) | (0.45) | (0.60) |
| Share of household-months | 32.4% | 67.6% | |

Each cell is $\theta$ from equation (4) estimated on rural or on urban households alone, with the units and windows of Tables 6 and 7, and the last column is the difference with its standard error. Residence is the survey's region type. Standard errors clustered by district in parentheses. *, ** and *** denote significance at the 10, 5 and 1 percent levels.

**Table A7:** All Coefficients of the Triple Difference for the Main Margins

| Outcome | $\beta$ $S$ | $\lambda$ $S \times$ NBFC | $\rho$ $S \times$ Post | $\mu$ NBFC $\times$ Post | $\theta$ $S \times$ NBFC $\times$ Post |
|---|---|---|---|---|---|
| Discretionary share (pp) | -0.62*** | 0.04 | 0.20 | -0.18 | -0.52* |
| | (0.17) | (0.21) | (0.22) | (0.24) | (0.31) |
| Protected index (%) | -0.7* | 0.3 | 0.7 | -0.8 | -3.2*** |
| | (0.4) | (0.5) | (0.5) | (0.6) | (0.8) |
| Log income, exposed occupations (×100) | -3.6*** | 0.2 | -0.2 | -0.5 | -0.2 |
| | (0.9) | (1.1) | (1.2) | (1.3) | (1.7) |
| Log total expenditure (×100) | -0.3 | 0.1 | -0.1 | -0.4 | -0.1 |
| | (0.5) | (0.6) | (0.7) | (0.7) | (1.0) |
| Non-bank borrowing (pp) | 0.24 | -0.07 | -0.05 | -0.46** | -0.66** |
| | (0.15) | (0.18) | (0.19) | (0.21) | (0.29) |
| Bank borrowing (pp) | 0.63* | -0.11 | 0.08 | 0.21 | -0.34 |
| | (0.38) | (0.43) | (0.44) | (0.48) | (0.69) |
| Moneylender borrowing (pp) | 0.41 | -0.12 | -0.08 | 0.26 | 0.64 |
| | (0.28) | (0.33) | (0.34) | (0.36) | (0.44) |
| Instalments paid on loans (%) | 1.2 | 0.9 | 0.6 | -3.4 | -6.7* |
| | (1.8) | (2.1) | (2.2) | (2.5) | (3.9) |

The five coefficients of equation (4) for each outcome, with the units of the tables in the text and log outcomes multiplied by one hundred, and $\hat{\beta}$ and $\hat{\rho}$ the averages of the state-specific coefficients weighted by district-blocks. The lower-exposure column of those tables is $\hat{\beta}+\hat{\rho}$ and the higher-exposure column is $\hat{\beta}+\hat{\rho}+\hat{\theta}$. $\hat{\lambda}$ is the pre-default difference in the weather response of the top tercile and $\hat{\mu}$ is the post-default change in the outcome in the top tercile without a shock. Standard errors clustered by district in parentheses. *, ** and *** denote significance at the 10, 5 and 1 percent levels.

**Table A8:** Persistence of the Credit Interaction

| | $\theta$ on the shock in block | | |
|---|---|---|---|
| Outcome | $t$ | $t-1$ | $t-2$ |
| Discretionary share (pp) | -0.52* | -0.25 | -0.07 |
| | (0.31) | (0.33) | (0.35) |
| Protected index (%) | -3.2*** | -1.7* | -0.5 |
| | (0.8) | (1.0) | (1.1) |
| Health (%) | -3.5*** | -2.1 | -0.7 |
| | (1.3) | (1.5) | (1.7) |
| Education (%) | -4.4** | -2.5 | -1.1 |
| | (1.8) | (2.0) | (2.1) |
| Milk and milk products (%) | -2.7*** | -0.9 | 0.2 |
| | (0.8) | (1.0) | (1.1) |
| Non-bank borrowing (pp) | -0.66** | -0.42 | -0.30 |
| | (0.29) | (0.31) | (0.32) |
| Moneylender borrowing (pp) | 0.64 | 0.37 | 0.17 |
| | (0.44) | (0.46) | (0.49) |

Each row is one estimate of equation (4) with the shock in the current block, the previous block and the block before that entered together, each with its own set of interactions, and the columns report $\theta$ on each. Units as in Tables 6 and 7. Standard errors clustered by district in parentheses. *, ** and *** denote significance at the 10, 5 and 1 percent levels.

**Table A9:** Alternative Exposure Measures

| Exposure measure | Discretionary share (pp) | Protected index (%) | Non-bank borrowing (pp) | Moneylender (pp) |
|---|---|---|---|---|
| Top tercile of the loan share (main) | -0.52* | -3.2*** | -0.66** | 0.64 |
| | (0.31) | (0.8) | (0.29) | (0.44) |
| Loan share, standardised (one s.d.), all districts | -0.25* | -1.5*** | -0.32** | 0.29 |
| | (0.15) | (0.5) | (0.14) | (0.23) |
| Top tercile of the rupee share | -0.48 | -2.9*** | -0.60* | 0.54 |
| | (0.32) | (1.0) | (0.31) | (0.46) |
| Top tercile of the share over 2015 and 2016 | -0.50 | -2.9*** | -0.62** | 0.57 |
| | (0.32) | (1.0) | (0.30) | (0.45) |
| Top tercile, gold loans included in the share | -0.50 | -3.0*** | -0.63** | 0.58 |
| | (0.31) | (1.0) | (0.30) | (0.45) |
| Top tercile, non-banks including fintech | -0.50 | -3.1*** | -0.63** | 0.60 |
| | (0.31) | (1.0) | (0.30) | (0.45) |
| Top tercile of the product-mix exposure (vehicle, property and housing by non-banks) | -0.46 | -2.8** | -0.56* | 0.50 |
| | (0.33) | (1.1) | (0.32) | (0.48) |
| Top tercile of the within-product share (loan-type mix held at national weights) | -0.43 | -2.5** | -0.54 | 0.47 |
| | (0.35) | (1.1) | (0.33) | (0.49) |
| Top tercile of the non-bank share of loans below Rs 50,000 | -0.55* | -3.4*** | -0.69** | 0.68 |
| | (0.32) | (1.0) | (0.31) | (0.46) |
| Top tercile, urban share interacted | -0.45 | -2.7** | -0.58* | 0.54 |
| | (0.35) | (1.1) | (0.32) | (0.46) |
| Top against middle tercile | -0.30 | -1.8* | -0.40 | 0.33 |
| | (0.32) | (1.0) | (0.30) | (0.45) |
| Middle against bottom tercile | -0.22 | -1.4 | -0.26 | 0.31 |
| | (0.31) | (1.0) | (0.29) | (0.43) |

Each cell is $\theta$ from equation (4) with the stated exposure measure in place of the top-tercile indicator, for the four margins in the column heads, with the units and windows of Tables 6 and 7. The standardised share gives the effect of one standard deviation of the June 2017 to May 2018 non-bank share of loans with all districts entered. The product-mix exposure is the share of the district's vehicle, property and housing loans originated by non-bank and housing-finance lenders over the same months, the within-product share averages the district's share within each product at national product weights, and the small-loan share uses loans below fifty thousand rupees alone. Standard errors clustered by district in parentheses. *, ** and *** denote significance at the 10, 5 and 1 percent levels.

**Table A10:** Alternative Index, Estimators, Samples and Designs

| Specification | Dry shock, $\beta^D$ (pp) | $\theta$ (pp) |
|---|---|---|
| Main: within-state design (state × shock and state × shock × post effects), one-month index, 2001 to 2021 thresholds | -0.68*** | -0.52* |
| | (0.21) | (0.31) |
| Pooled design, weather response and its change at the default common across states | -0.68*** | -0.64** |
| | (0.21) | (0.26) |
| Thresholds from the full 1901 to 2024 record | -0.63*** | -0.48 |
| | (0.22) | (0.32) |
| Three-month accumulation window | -0.84*** | -0.57 |
| | (0.28) | (0.45) |
| Six-month accumulation window | -0.94** | -0.45 |
| | (0.41) | (0.58) |
| Index winsorised at ±3 | -0.66*** | -0.51* |
| | (0.21) | (0.31) |
| Thresholds computed within calendar month | -0.66*** | -0.50 |
| | (0.21) | (0.31) |
| Monsoon months only | -0.92*** | -0.71 |
| | (0.34) | (0.52) |
| Cell-level intensity share | -0.97*** | -0.69 |
| | (0.33) | (0.49) |
| Count of dry months in the block (per month) | -0.45*** | -0.34* |
| | (0.14) | (0.20) |
| Straddling blocks coded by the majority of their months | -0.69*** | -0.50 |
| | (0.21) | (0.31) |
| Blocks with a month in each tail dropped | -0.70*** | -0.51 |
| | (0.23) | (0.33) |
| District covariates of Table A3 × shock × post | -0.65*** | -0.40 |
| | (0.22) | (0.36) |
| Farm-loan-waiver states excluded | -0.72*** | -0.56 |
| | (0.24) | (0.37) |
| Farming households excluded | -0.57** | -0.47 |
| | (0.23) | (0.35) |
| Survey weights applied | -0.71*** | -0.54 |
| | (0.24) | (0.33) |
| Pre-period block −3 dropped | -0.68*** | -0.50 |
| | (0.21) | (0.32) |
| de Chaisemartin and D'Haultfœuille estimator | -0.59** | -0.42 |
| | (0.24) | (0.35) |
| Districts in the panel since 2014 | -0.71*** | -0.54 |
| | (0.24) | (0.35) |
| Blocks containing November 2016 dropped | -0.70*** | -0.53* |
| | (0.22) | (0.32) |
| Wet tail in place of the dry tail | -0.59*** | -0.17 |
| | (0.21) | (0.32) |
| Monthly frequency | -0.19 | -0.10 |
| | (0.16) | (0.25) |

The first column is $\beta^D$ from equation (3) and the second is $\theta$ from equation (4), two separate regressions per row, both for the discretionary budget share in percentage points, under the stated alternative. The waiver-state row excludes Madhya Pradesh, Chhattisgarh, Rajasthan and Karnataka, and the farming-household row excludes small and organised farmers. Standard errors clustered by district in parentheses. *, ** and *** denote significance at the 10, 5 and 1 percent levels from the district-clustered errors in every row.

**Table A11:** Inference for the Main Estimate

| Method, $\theta = -0.52$ on the discretionary share | Standard error | $p$-value |
|---|---|---|
| District clusters (518), conventional | 0.31 | 0.094 |
| Spatial (Conley) standard errors, 200 km cutoff | 0.33 | 0.115 |
| Spatial (Conley) standard errors, 500 km cutoff | 0.35 | 0.137 |
| Homogeneous-region clusters (102) | 0.34 | 0.126 |
| State clusters (27), wild cluster bootstrap | 0.36 | 0.148 |
| Exposure terciles permuted across districts within state (re-sorts only the ten states with both terciles), 1,000 draws | | 0.108 |
| Exposure terciles permuted across all districts, 1,000 draws | | 0.083 |
| Whole district weather series permuted, 1,000 draws (sharp null of no weather effect) | | 0.121 |
| Pooled design, weather response common across states, $\theta = -0.64$: district clusters | 0.26 | 0.014 |

The main $\theta$ on the discretionary share, $-0.52$, under alternative treatments of the error structure. Spatial standard errors follow Conley (1999) with a uniform kernel and the stated cutoff between district centroids. The homogeneous regions are the survey's 102 sampling regions. The state row uses the wild cluster bootstrap of Cameron et al. (2008) with 999 draws. The exposure permutation reassigns the tercile indicators across districts within state and re-estimates $\theta$ 1,000 times, as Breza and Kinnan (2021) permute exposure across districts, and the $p$-value is the share of draws at least as large in absolute value, and the second permutation reassigns them across all districts. The weather permutation reassigns whole district weather series across districts and tests the sharp null that weather has no effect on the outcome, and the last row reports the pooled design, which uses the cross-state variation as well, under district clusters.

**Table A12:** Falsification Tests

| Test | Estimate | Std. error |
|---|---|---|
| Salaried households, income, dry tail | -0.0021 | (0.0068) |
| Salaried households, income, wet tail | 0.0027 | (0.0066) |
| Retired households, income, dry tail | 0.0009 | (0.0079) |
| Retired households, income, wet tail | -0.0148* | (0.0076) |
| Retired households, income, wet tail, member count unchanged | -0.0091 | (0.0080) |
| Salaried households, discretionary share, dry tail | -0.061 | (0.184) |
| Retired households, discretionary share, dry tail | -0.088 | (0.191) |
| Salaried households, $\theta$ on non-bank borrowing | 0.07 | (0.37) |
| Retired households, $\theta$ on the discretionary share | -0.07 | (0.39) |
| Next year's shock on this year's discretionary share | 0.074 | (0.198) |
| District originations of Rs 50,000 or more, dry tail (log) | 0.0018 | (0.0043) |
| $\theta$ with the default dated September 2017, sample ending August 2018 (pp) | 0.07 | (0.37) |
| Exit from the clean panel next wave, dry tail (pp) | -0.59* | (0.35) |
| Exit from the clean panel next wave, $\theta$ (pp) | -0.18 | (0.57) |

Each row is a separate regression using the same specification as the corresponding main result, restricted to the stated group or sample. The originations row regresses log district originations of loans flagged at fifty thousand rupees or more on the dry-tail indicator with district and month fixed effects in the bureau panel. The two exit rows use an indicator for a household present in the clean panel in one wave and absent in the next, whether not interviewed or interviewed with an incomplete or flagged block, so the base exceeds the share of households the survey loses between waves, over waves 7 to 17, first on the dry-tail indicator and then in equation (4). Standard errors clustered by district in parentheses. *, ** and *** denote significance at the 10, 5 and 1 percent levels.

**Table A13:** Benchmarks from the Literature

| Study and setting | Variation | Household effect | Lending effect, scaling |
|---|---|---|---|
| Breza and Kinnan (2021), India, 2010 microfinance ordinance | Removal: district exposure, one s.d. | Monthly consumption −151 Rs (2.7% of mean) and casual wage −2.2% | Microcredit per household −626 Rs |
| Jack and Suri (2014), Kenya, mobile money | Addition: M-PESA use | Shock lowers non-users' consumption by about 7% and leaves users unaffected | |
| Suri, Bharadwaj and Jack (2021), Kenya, digital loans | Addition: credit-score cutoff | 6.3 pp less likely to forgo expenses after a shock | |
| Cramer et al. (2025), India, same bureau, IL&FS | Removal: non-bank funding | Non-bank lending response to a weather shock −1.7 log points after the default, fintech unchanged | No household outcome |
| This paper, India, IL&FS, household panel | Removal: top against bottom tercile | Discretionary share −0.52 pp, protected spending −3.2%, non-bank incidence −0.66 pp | Non-bank rupees −12% (−0.13 log points), and no rupee scaling, since the survey records no loan amounts |

Published estimates as reported in each paper, at the pages or tables given in the text. Our own entries are from Tables 1, 5, 6 and 7.

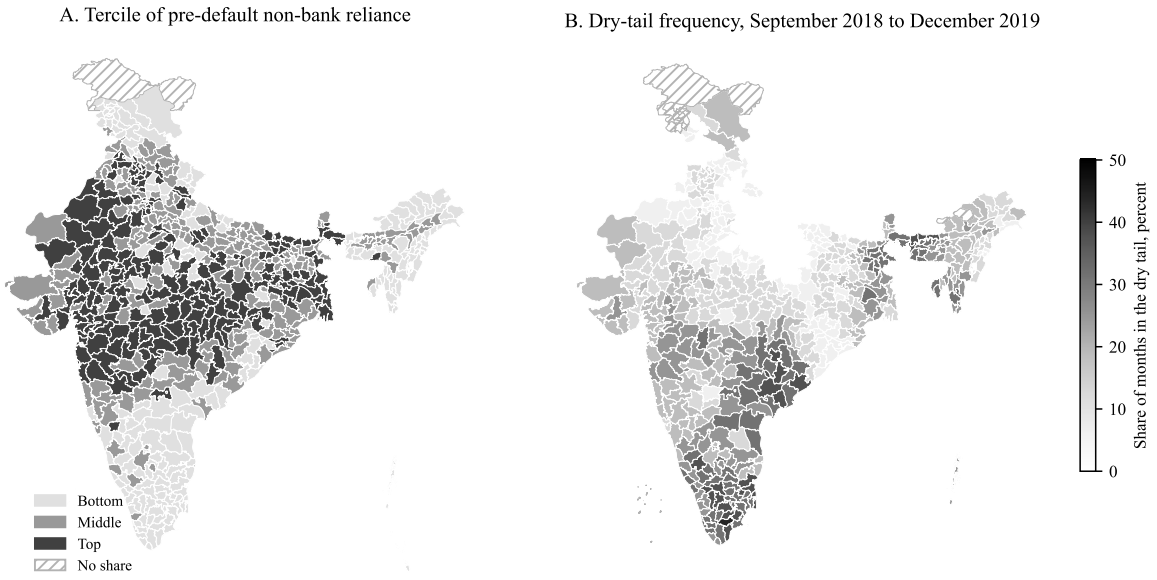



**Figure A1:** Exposure Terciles and Post-Default Dry-Month Frequency by District
Panel A shades each Census-2011 district by its tercile of the non-bank and housing-finance share of loans originated from June 2017 to May 2018, with the territory outside the 640 bureau districts hatched. Panel B shades each district by the share of its months from September 2018 to December 2019 in the dry tail of its own water-balance history, with the 18 districts without an index hatched.

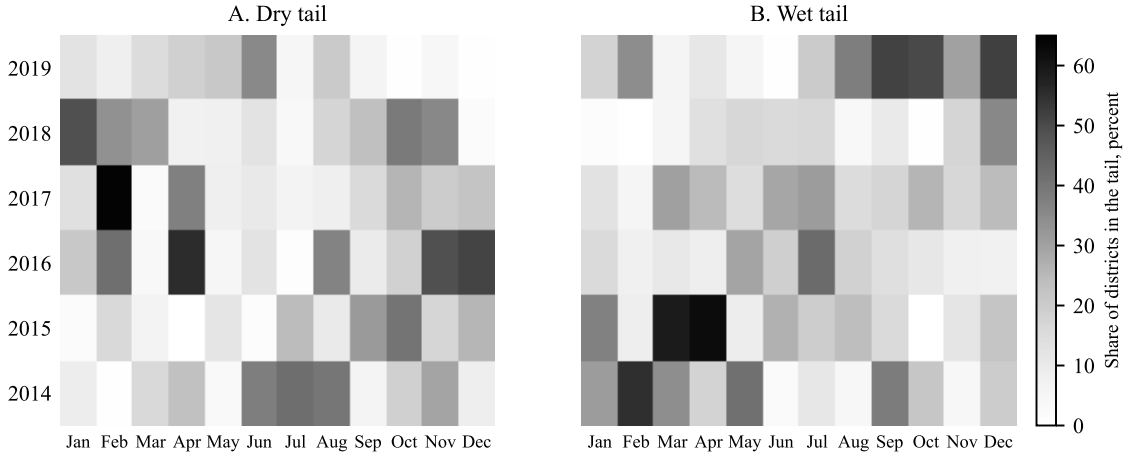



**Figure A2:** Incidence of Dry and Wet Shocks by Calendar Month and Year
Each cell is the share of the 623 districts with an index whose water balance fell in the dry tail (Panel A) or the wet tail (Panel B) in that month, under the thresholds of equation (2).

# B Additional Robustness Checks

Section 4.3 summarises the checks on the exposure measure, the weather index, the design and inference. This appendix records each choice along with what the alternative does, and gives

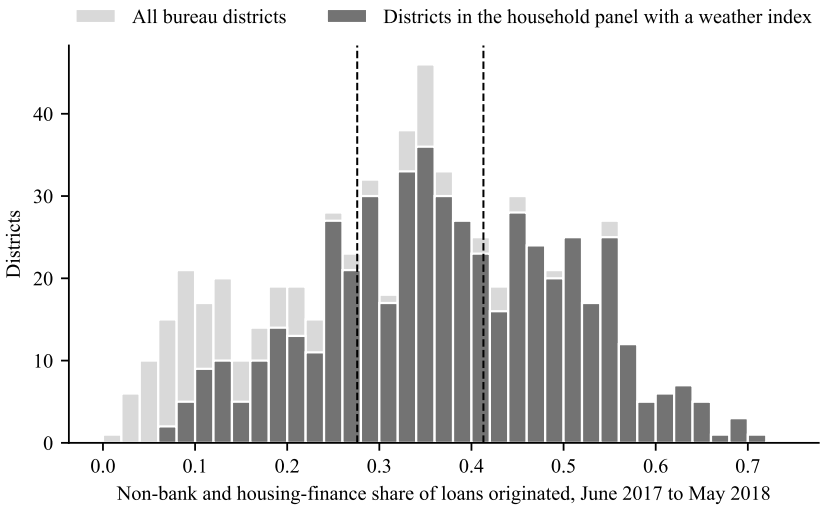



**Figure A3:** The Distribution of Pre-Default Non-Bank Reliance across Districts

The non-bank and housing-finance share of loans originated from June 2017 to May 2018 in each of the 640 bureau districts, with non-bank gold loans left out, and in darker shading the 518 districts that enter the household estimates. The dashed lines are the tercile cutoffs, 0.28 and 0.41.

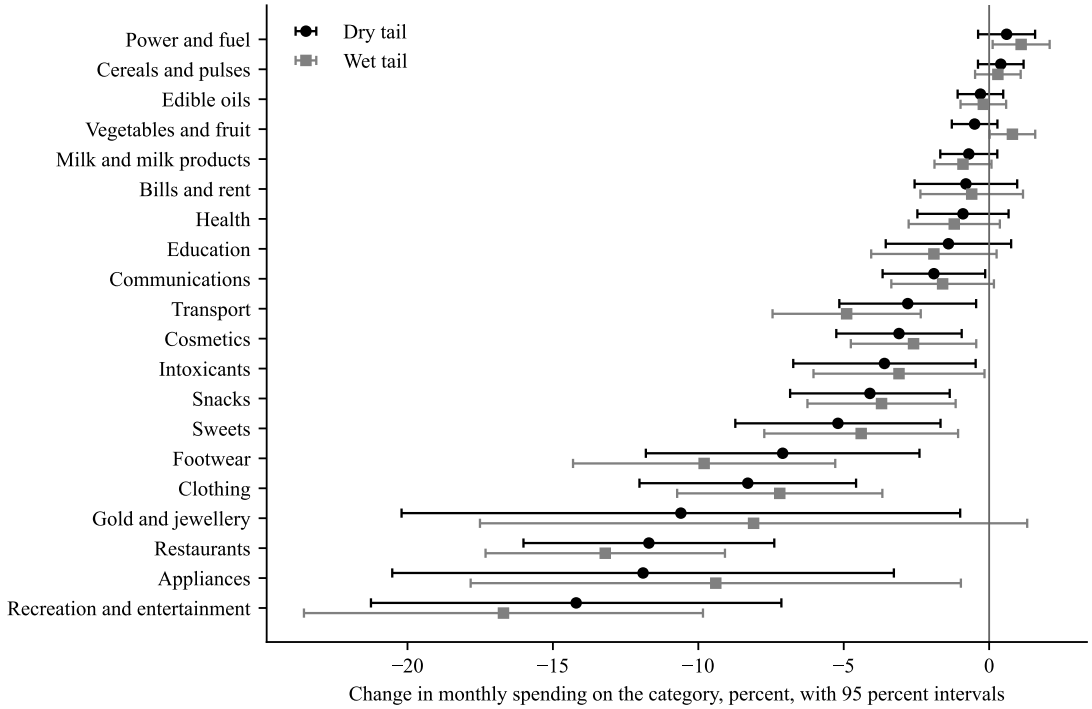



**Figure A4:** Expenditure Responses by Category, Both Tails, All Households

Estimates of $\beta^D$ and $\beta^W$ from equation (3) with the inverse hyperbolic sine of rupees spent on each category as the outcome, read as percent changes, with 95 percent intervals from standard errors clustered by district, over 2,196,418 household-blocks. Categories are ordered by the dry-tail response, from the least postponable at the top to the most at the bottom, and Figure 3 in the text keeps this order.

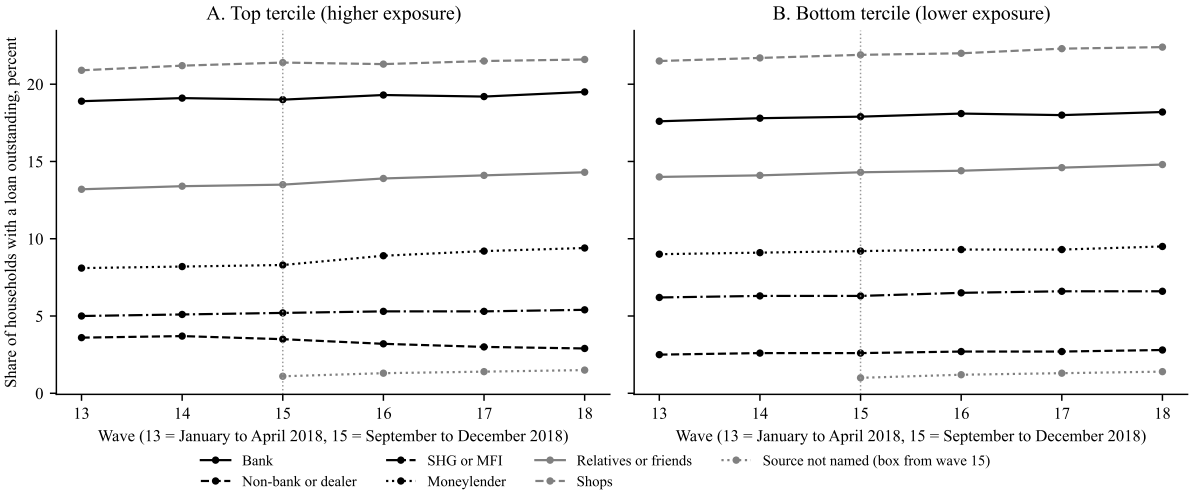



**Figure A5:** Borrowing Incidence by Source and Exposure Tercile, Waves 13 to 18
The share of households reporting a loan outstanding from each source in each wave, for the top and bottom terciles of frozen non-bank reliance. The dotted line marks wave 15, September to December 2018, which contains the default and in which the questionnaire added a box for a loan whose source the household cannot name. The named sources show no break at the wave in either tercile.

the details of the inference procedures.

*The Exposure Measure.* Appendix Table A9 enters exposure as the share itself standardised to unit variance, which gives the effect of a one standard deviation difference in reliance, about half the tercile contrast on every margin, since the top and bottom terciles sit 2.3 standard deviations apart. Building the share from rupees in place of loans, over 2015 and 2016 in place of the twelve months before the default, with gold loans kept in it, with financial-technology lenders counted among the non-banks, or from the vehicle, property and housing lending of non-banks alone, the products whose lenders depended most on commercial paper, moves the estimates by less than a standard error. Two exposures answer the objection that the share mostly measures the local product mix. A share built within each product and averaged with national product weights, so that a district's loan mix cannot drive it, gives $-0.43$ with a standard error of 0.35 on the discretionary share, and a share built on loans below fifty thousand rupees alone, which are the loans the exposed occupations hold, gives $-0.55$ with a standard error of 0.32, slightly above the main estimate. Under the 2015 to 2016 share the event study of Figure 2 keeps its shape, with the coefficient three blocks before the default at 0.33 with a standard error of 0.41 and a joint pre-period $p$-value of 0.38, so the largest pre-period coefficient is not an artefact of building the share from lending in that block. Dropping wave 15, the wave that contains the default and the questionnaire change, from the borrowing margins leaves the non-bank $\theta$ at $-0.63$ with a standard error of 0.30 and the moneylender $\theta$ at 0.60 with a standard error of 0.46, and counting the unnamed-source box as a source, informal first and then formal, leaves the non-bank $\theta$ at $-0.64$ with a standard error of 0.29 and the transition from formal credit alone to no source at 3.5 points with a standard error of 2.3.

*Logs with Zeros.* The category rows enter rupees through the inverse hyperbolic sine, and Chen and Roth (2024) show that the coefficient of such a transformation depends on the unit of the outcome whenever the treatment moves households between zero and positive spending, which it does for education, bought in 56 percent of the household-blocks. Poisson pseudo-maximum likelihood with household and state-by-block fixed effects is unit-free, and on rupees spent it gives a $\theta$ of $-3.2$ percent for health with a standard error of 1.2, $-3.9$ for education with a standard error of 1.7, $-2.6$ for milk with a standard error of 0.8, and $-6.2$ for instalments with a standard error of 3.7, against $-3.5$, $-4.4$, $-2.7$ and $-6.7$ in Table 7, so the reading of those rows does not depend on the transformation.

*The Estimator for Switching Treatments.* The row of Appendix Table A10 labelled with de Chaisemartin and D'Haultfœuille (2020) applies their estimator for a binary treatment that switches on and off to district-block means of the outcome within each of the four cells of the design, the two terciles before and after the default, so that each cell yields the effect of a dry month on the districts that switch into it against those that stay, and $\theta$ is the double difference of the four with a standard error from a district-clustered bootstrap. The weights the two-way fixed-effects regression places on the switching cells are all positive in these data.

*The Accumulation Window.* Equation (8) accumulates the water balance over $k$ months, and $k$ is a substantive choice. A one-month window measures meteorological drought, meaning a dry month, and fires in roughly a fifth of months by construction. A three-month window corresponds to a crop season and is the conventional choice for agricultural drought, and six and twelve month windows measure soil moisture and groundwater. Our main specification uses $k = 1$, matching Cramer et al. (2025) so the estimates are comparable to theirs, and we examine $k = 3$ and $k = 6$ as robustness. The longer windows are more persistent, and the persistence has an inference cost as well as a measurement benefit, since the district-clustered standard errors widen from 0.31 for $\theta$ under the one-month index to 0.45 under three months and 0.58 under six, and the spatial standard errors widen by more for the longer windows than for the main index. The weather coefficients under the longer windows are larger, which is the measurement benefit, and the credit interactions are not distinguishable from zero under either longer window, which is the cost.

*The Baseline Period.* The quantiles $q_d(p)$ in equation (2) are computed over January 2001 to December 2021, the window Cramer et al. (2025) use, which is also the recommendation of our own validation of the index, since thresholds taken over a record that begins in 1901 compare the study years with a climate that no longer holds. Over the full record instead, the mean twentieth and eightieth percentiles sit 0.09 and 0.08 index points higher, since the recent decades are drier than the century, so the study-period shock rates move from 19.0 and 20.4 percent to 21.4 and 18.4 percent, and Appendix Table A10 shows that the estimates move by a tenth.

*Spatial Weighting.* Equation (10) weights cells by population, and the alternative is to weight

by the area of the intersection, which produces a series correlated 0.9992 with the population-weighted one, with a mean absolute difference of 0.021 index points. We prefer population weighting on the argument that the shock should be measured where people are rather than where land is, and the data say the choice is close to immaterial.

*An Intensity-Based Alternative.* Thresholding a district mean reintroduces the spatial averaging that Appendix D identifies as one reason the wet tail is poorly measured. An alternative that avoids it is to threshold at the cell level first and then aggregate, giving the population share of the district already in deficit,

$$\text{dryshare}_{dm} = \sum_{c} \omega_{cd}\, \mathbf{1}[\text{SPEI}_{cm} < q_c(0.20)]\,. \tag{6}$$

This is continuous, not binary, and immune to the averaging problem. We report the binary form in the paper because it is the form the comparison literature uses, and the intensity form as a robustness check.

*Seasonal Imbalance in Threshold Crossings.* The quantiles in equation (2) are computed pooling across calendar months, following the way the underlying grid was built. A twentieth-percentile cut ought then to fire in about 20 percent of months in every month, and it does not. Observed dry-tail rates run from 28.7 percent in December and 25.1 percent in February down to 15.0 percent in May and 16.3 percent in September. The cause is the standardisation in equation (9), since in months when normal rainfall is near zero the fitted distribution is concentrated close to zero, so a small absolute departure maps to an extreme quantile and the index takes implausibly large values, with 113 cell-months in the record above four in absolute value and eighteen above five, reaching $-8.04$, clustered in the dry season. This is measurement error in the treatment, and it attenuates. Percentiles can be computed within calendar month, the index can be winsorised at plus or minus three, or the sample can be restricted to monsoon months, which has direct precedent in Jayachandran (2006). Appendix Table A10 reports all three, the monsoon restriction at the cost of roughly two thirds of the sample. Eleven percent of district-blocks hold a month in each tail and enter both indicators, and dropping them moves the estimates by less than a standard error.

*Monthly Frequency.* At monthly frequency the discretionary-share estimate is $-0.19$ with a standard error of 0.16, against $-0.68$ at four-month blocks (Appendix Table A10). The two are not two readings of the same thing. Spending in the survey is a block average by construction, so a monthly shock indicator is a noisy proxy for the quantity the block design measures, whether any month of the block was dry, and the attenuation to between a quarter and a third is close to what that proxy predicts when a treated block holds one and a half dry months in four on average. The block estimate is the one the instrument supports, and the monthly estimate is reported so that a reader can see the cost of ignoring the instrument's structure.

*Inference Procedures.* The spatial standard errors in Appendix Table A11 apply the estimator of Conley (1999) to the household-block residuals aggregated to district by block, with a uniform kernel and cutoffs of 200 and 500 kilometres between district centroids taken from the Census-2011 district polygons. The wild cluster bootstrap clusters by state, imposes the null, and draws 999 Rademacher weights (Cameron et al., 2008). The exposure permutation holds every district's weather series and outcomes fixed, reassigns the tercile indicators at random across the districts of each state 1,000 times, and re-estimates $\theta$ in each draw, which is the procedure Breza and Kinnan (2021) use to test their exposure measure, and the $p$-value is the share of draws in which the absolute estimate is at least the actual one. Within state, the reassignment changes nothing in the seventeen states whose districts all sit in one tercile, so the test re-sorts only the ten mixed states and leaves the cross-state variation in place, which is why Appendix Table A11 also permutes the terciles across all districts, a test that breaks the spatial contiguity of the true terciles and so has thinner tails than the truth, and reports the pooled design's district-clustered error for comparison. The weather permutation reassigns entire district weather series to other districts, preserving the serial and seasonal structure of each series while breaking its link to the place, and tests the sharp null that weather has no effect on the outcome, which is a stronger null than $\theta = 0$ and is reported for completeness. The sharpened $q$-values of Anderson (2008) are computed across the twenty rows of Tables 6 and 7 together, the secondary margins, from the district-clustered $p$-values.

# C Data Appendix

*Bureau Lender and Product Taxonomy.* The bureau extract classifies lenders as public banks, private banks, non-bank and housing-finance companies, financial-technology non-banks, and others, and products into thirteen categories, mortgage, vehicle, personal, agriculture, gold, microfinance, business, credit card, consumer durable, education, mudra, farm credit card and others. The mortgages of the non-bank group are mostly housing-finance lending, since housing-finance companies take about seven in ten of the group's mortgage inquiries, and the rest of its book is non-bank lending. Each origination cell also carries the borrower's credit-score band, a flag for a prior relationship with the lender, the borrower's gender, and a flag for a sanctioned amount of fifty thousand rupees or more, with the amount unknown for a small number of loans, and we use none of the first three, since the prior flag has documented breaks and the score bands begin a few months into the panel. Two series are left out of every total and of the exposure share, because their reporting breaks inside the window. Gold loans by non-bank lenders fall by four fifths in August 2017 and return in April 2019 with a different borrower mix, and the farm credit cards reported by lenders outside the four main groups jump from about twenty thousand a month to four hundred and forty thousand in February 2017, nine tenths of them in one state. Loans with no reported amount, under one percent of loans a year, and loans recorded at the

bureau's stand-in value of a billion rupees, count as loans and add no rupees, and the handful of cells that still average over a hundred million rupees per loan are carried in the totals and left out of every series that uses rupees per loan. Asset reconstruction companies re-report a small number of non-bank loans sold to them, under one in sixty loans of the residual lender group, and those twins are removed. Whether a loan sold by a stressed non-bank to a bank is recorded under the seller or the buyer is not documented, so if such sales were re-reported the bank series would be overstated after the default, which works against the finding that bank rupees did not rise. The rows the bureau removed before delivery, loans without a score band, loans with negative amounts and loans in pincodes outside the national list, are of unknown number. The exposure share of equation (1) counts non-bank and housing-finance lenders in the numerator and all lenders in the denominator, over originations of all other products from June 2017 to May 2018, and the product-mix exposure of Appendix Table A9 does the same over vehicle and mortgage lending alone, the bureau's mortgage category holding both housing loans and loans against property. 12 districts in which more than a twentieth of credit inquiries over 2015 to 2019 fall in pincodes that never record a loan are dropped from the lending regressions. The survey's own indicator for holding a farm credit card is the one household-side credit instrument named identically on both sides, and we do not use it, since the bureau's farm credit card series is the one with the 2017 break. The extract begins a year after the household panel, so the lending regressions and the comparison with transfers, which use the bureau month by month, run from 2015, while results drawn from the household data alone use the full panel from January 2014.

*Interview Blocks and Their Coding.* The unit of observation is the interview block, the four calendar months that one interview reports on. The survey records the reference months of each interview, so the block is observed directly. A block is complete when all four months are reported, and the estimates use complete blocks. Each block is assigned to the state-by-block effect of its reference months and to a weather treatment that is one if the dry tail fires in any of its four months, with the count of such months used as an alternative. The post-default indicator is one for blocks whose four reference months all fall in September 2018 or later and zero for blocks whose months all fall before, and the blocks whose months straddle the default are dropped in the main estimates and coded by the majority of their months in a check. The wave-level modules, borrowing, assets, gold saving and employment, are recorded on the date of the interview, one to four months after the reference months, and are assigned to the wave in which the interview falls, so for them the post-default indicator, the event-study index and the state-by-block effect follow the interview wave, with waves 15 to 18, interviewed from September 2018, as the post-default waves, while the weather shock is the dry tail over the four reference months the interview reports on. A block whose reference months straddle two calendar blocks is assigned to the calendar block that holds three or more of them, and a block split two and two, which the interviews of March, July and November produce, to the calendar block in which the interview falls, so that in the majority-coded check the July to October 2018 window counts

as post-default. The panel's reference months end in December 2019 and its waves with wave 18, interviewed September to December 2019, so the final calendar block of reference months holds fewer complete windows than the others, which is one reason its event-study coefficient is the least precise. Log income is undefined for the 0.3 percent of complete blocks that report no income, which drop out of the log-income rows, the wage-income rows use the 62 percent of blocks with wage income, and neither share responds to the dry tail, by 0.01 and 0.1 percentage points with standard errors of 0.02 and 0.2. Between two and three fifths of early blocks report the same income in all four months, so the median month-on-month change in income of 3 percent in Section 3 is an average over blocks with no change and blocks with some, the median coefficient of variation of spending within a block is under 0.10, and the demeaned means by recall position alternate instead of decaying, which is the recall structure that the block design cancels. Category spending enters the regressions as the inverse hyperbolic sine of rupees, so household-blocks with no purchase of a category stay in the sample. The cleaned panel is the set of households with at least one complete block after the survey's own response flags are applied. In the Diwali months of the panel the fifteen heads fall about one percent short of the total, a festival artefact that the state-by-block effects absorb.

*Exposure Groups.* Households are assigned to one of ten groups from the occupation of the household head at first observation and the industry of the head's work. The survey's occupation classes supply agricultural labourers, small farmers, organised farmers, salaried employees and retired households directly. Wage labourers are split by industry into construction labourers, whose industry is outdoor work, construction, mining, brick and cement, fishing, forestry or cultivation, and indoor casual labourers in every other industry. The self-employed are split the same way into the outdoor self-employed, in construction, animal husbandry, fishing and forestry, and the other self-employed, which holds traders, shopkeepers, hawkers and the self-employed in manufacturing and services. The residual group holds households whose head is not in the labour force, home makers, students and the unoccupied among them. Appendix Table A1 reports the size, rural share, income, borrowing incidence and budget shares of each group. Agricultural labourers and small farmers are the most rural, agricultural labourers and construction labourers the poorest, and construction labourers the largest exposed group, at 28,912 households.

*The Consumption Basket.* The survey's spending items are hierarchical, so summing all of them roughly doubles total spending, and fifteen heads partition total monthly expenditure, food, intoxicants, restaurants, entertainment, clothing and footwear, cosmetics and toiletries, bills and rent, instalments, appliances, power and fuel, transport, communication, education, health and miscellaneous items. The twenty categories of Appendix Table A4 are taken from these heads and their leaves, with the essential or discretionary label assigned in advance of estimation. On this basis cereals and pulses, edible oils, vegetables and fruit, milk, health, power and fuel, and bills and rent are essential, and restaurants, recreation, clothing, footwear, toiletries

and cosmetics, gold and jewellery, intoxicants, communications, transport, education, appliances, sweets and snacks are discretionary. Electricity is not postponable for a household with a metered connection even though richer households spend more on it, which is why the classification follows deferability and not the Engel elasticity where the two diverge. The recreation category is the survey's recreation aggregate, which contains the entertainment head along with games and toys and electronic storage, so entertainment is not counted separately, and the residual is what the fifteen heads contain that the twenty categories do not. The milk category is the survey's milk and milk products head, and the gold and jewellery category is its gems and jewellery head, most of which is gold. Category responses in Figures A4 and 3 are coefficients on the inverse hyperbolic sine of rupees spent on the category, read as percent changes, the protected index of Table 7 is the log of rupees spent on health, education and milk together, and the shares in Tables 4 and 5 are the category's rupees as a share of total monthly expenditure, in percentage points.

*Borrowing Sources, Purposes and States.* The borrowing module records, for each wave, whether the household holds a loan outstanding from each of eleven sources, banks, moneylenders, employers, relatives and friends, non-bank finance companies or dealers, self-help groups, microfinance lenders, chit funds, credit cards, shops and other sources, and from wave 15 a twelfth box for a source the household cannot name, and whether it has borrowed for each of eleven purposes, housing, education, medical expenses, weddings, consumption, consumer durables, business, investment, the repayment of other debt, vehicles and other purposes. The module's sources were added through the first three years of the panel, with shops entering in 2016, and reported incidence rises from 6 percent of households in 2014 to 49 percent in 2019, roughly half of which is the instrument, which is why every borrowing estimate uses waves 13 to 18. The formal sources are banks, non-bank finance companies or dealers, self-help groups, microfinance lenders and credit cards, and the informal sources are moneylenders, relatives and friends, shops, employers, chit funds and other sources. Where self-help groups belong is arguable, since they lend their members' savings and bank credit, and placing them with the informal sources moves the formal-only share of waves by three points and no transition estimate by more than a standard error. The unnamed-source box enters neither list, so a household that reports only that box is in the state with no source. The four borrowing states of Table 8 are defined from the formal and informal indicators of the same wave, and a transition is a change of state between two consecutive waves in which the household was interviewed, with no wave missing between them, so households in the strata that the survey does not visit every wave contribute fewer pairs. The transition estimates condition on the origin state, so that a change in the share of households holding formal credit, which the contraction itself produces, cannot move them, and a pair takes its post-default indicator, weather shock and state-by-block effect from the later interview. The survey retains 88 percent of the households interviewed in the wave of the default a year later, so a household contributes most of its five possible pairs. Three

borrowing purposes have no bureau product code, medical expenses, marriage, and borrowing to repay other debt, the last being the classic distress signature, invisible to the bureau because no lender labels a loan as debt consolidation. Income and spending are not independent measurements, because the interviewer shows the household its implied saving and offers it the chance to revise both before the interview closes, so the income losses and the spending responses are reconciled numbers and should be read together.

*Transfers.* Receipt of government transfers and of private transfers is recorded as income from each source in the income module. The government transfer variable excludes wages from the rural employment guarantee, which the survey records as labour income, and its incidence rises from 1.1 percent of household-months in 2014 to 60.6 percent in 2019, a rise that is largely real and partly shares the shape of two other series whose collection changed in 2017 and 2018. For the comparison of Appendix E each buffer is a district-level measure frozen before the panel, the share of the district's surveyed households that received each transfer over the first two waves, interviewed in 2014, leaving each household's own receipt out, and the district's bureau originations per person over 2015 and 2016, each standardised to unit variance.

*Sampling Frame, Weights and Coverage.* The panel is contested as a measure of India. Somanchi (2021) shows that it misses the poor at a scale that invalidates level estimates, with illiterate adults at one to two percent of the sample against sixteen to twenty-eight percent in the National Family Health Survey, Drèze and Somanchi (2024) argue that reweighting cannot repair the omission, and Sinha Roy and van der Weide (2025) show that the 2017 expansion of the sampling frame brought in areas whose households had eighteen percent lower per-capita consumption in the 2011 census. Abraham and Shrivastava (2022) find the bias largely non-differential across categories, and within-household and between-group estimands survive non-differential bias, so no poverty rate, participation rate or consumption level appears anywhere in the paper, and we make no claim about women's labour supply, where the survey is known to be badly wrong. The missing households are also those with the least access to formal credit and the fewest alternative margins, so if credit access changes the coping portfolio most for the poorest, every estimate reported here is attenuated relative to the population effect. The survey supplies household weights for population statements, and the regressions do not use them, because their estimands are changes within households and the weights would reweight the surveyed population towards a population the survey does not represent well, and Appendix Table A10 reports the main estimate with the weights applied. District coverage is not constant, since a few strata, five northeastern states and a handful of districts elsewhere, are not surveyed in every wave, and the household fixed effects absorb their entry and exit. A hundred of the survey's districts hold three or fewer primary sampling units, which is why Appendix Table A11 also clusters by the survey's homogeneous regions.

*District Crosswalk and Sample Restrictions.* Bureau records at pincode level are aggregated to Census-2011 districts with a population-weighted crosswalk from pincodes to districts, and the

fraction of loans placed by a fill-in rule or a name match rather than a direct pincode match is below one percent nationally. The weather index is aggregated from half-degree cells to the same districts with 2011 gridded population weights, as Appendix D describes. 18 of the 641 district codes have no index, 14 of them in Jammu and Kashmir, because SPEIbase masks the disputed territory, along with two districts of Arunachal Pradesh, Lakshadweep and one unadministered code, and they are excluded throughout. The survey's households are placed in Census-2011 districts through a crosswalk from the survey's own district list, and three sampling units in Delhi that span several Census districts are split across them in proportion to population. The household panel expanded its sampling frame in 2017, and Section 4.3 reports that restricting the sample to the districts present since 2014 leaves the estimates materially unchanged.

# D Construction and Validation of the Weather Index

Section 3 defines the treatment, and this appendix gives the construction in full and reports the validation behind the claim that the two tails are not equally well measured.

*The Index.* For a grid cell $c$ in month $m$ the climatic water balance is

$$D_{cm} = P_{cm} - \mathrm{PET}_{cm} \tag{7}$$

where $P$ is precipitation and PET is potential evapotranspiration computed by the Penman-Monteith method. Because a crop does not fail on one dry day, the balance is accumulated over a window of $k$ months,

$$D^{k}_{cm} = \sum_{j=0}^{k-1} D_{c,m-j}. \tag{8}$$

The accumulated balance is fitted to a three-parameter log-logistic distribution separately for each cell and each calendar month, and the index is the standard normal quantile of the fitted cumulative distribution,

$$\mathrm{SPEI}^{k}_{cm} = \Phi^{-1}\left(F_{c,\,\mathrm{mo}(m)}\left(D^{k}_{cm}\right)\right). \tag{9}$$

Fitting within calendar month makes a dry December comparable to a dry July, since each is judged against the distribution of Decembers or of Julys in that same cell. SPEIbase v2.9 is built on CRU TS 4.07 precipitation and runs from 1901 to 2024.

*Aggregation to Districts.* The index is produced on a grid and the analysis needs districts, and the two do not line up. The median Indian district spans roughly 2.9 effective half-degree cells with only about 47 percent of its area in its largest one, so collapsing the grid to the administrative unit is a real choice. We aggregate using population weights,

$$\mathrm{SPEI}_{dm} = \sum_{c} \omega_{cd}\,\mathrm{SPEI}_{cm}, \qquad \omega_{cd} = \frac{\mathrm{pop}_{cd}}{\sum_{c'} \mathrm{pop}_{c'd}} \tag{10}$$

where $\text{pop}_{cd}$ is the population living in the intersection of cell $c$ and district $d$, taken from the WorldPop 2011 raster, which reproduces the 2011 census total to within a factor of 1.0003. The percentile thresholds in equation (2) are then recomputed on the aggregated district series and not inherited from the constituent cells. The choice matters, since an average of three cells has lower variance than any one of them, so cell thresholds applied to a district mean would systematically under-flag districts.

*Validation Against Documented Events.* We validated the index against documented events using a rank test that cannot be tuned by choosing a flattering month. For each event we ask where that month ranks for extremeness among all 243 months from January 2000 to March 2020 in the region actually affected. A real drought should rank near the top of the driest months and a real flood near the top of the wettest.

**Table A14:** Rank Test of the Weather Index against Documented Events

| Documented event | Region | SPEI | Rank | Detected? |
|---|---|---|---|---|
| All-India drought, Jul 2002 | Nationwide | -1.52 | 3 / 243 | Yes |
| All-India drought, Jun 2009 | Nationwide | -1.11 | 8 / 243 | Yes |
| Chennai floods, Dec 2015 | Tamil Nadu | 0.90 | 36 / 243 | No |
| Mumbai floods, Jul 2005 | Maharashtra | 0.79 | 37 / 243 | No |
| Kerala floods, Aug 2018 | Kerala | 0.82 | 47 / 243 | No |
| Uttarakhand floods, Jun 2013 | Uttarakhand | 1.05 | 48 / 243 | No |

Rank of the documented event month among all 243 months from January 2000 to March 2020, by mean index value in the affected region. Regions are built from district polygons joined to states. An event counts as detected if it falls in the ten most extreme months of the 243.

The drought side passes cleanly, since the 2002 all-India drought is the third driest month in the two decades and the 2009 drought is the eighth. Extending the test from two hand-picked events to all twenty-one monsoon seasons, the four driest seasons by June to September mean index are 2002, 2004, 2009 and 2014, and all four are years the India Meteorological Department classifies as deficient. The index misses 2015, which the Department also classifies as deficient.

The wet side does not pass, because none of the four major floods ranks in the top thirty months for its region, and all four sit between the eightieth and eighty-sixth percentiles. Three mechanisms compound to produce this pattern, the first of which is that monsoon rainfall is convective and spatially patchy, so within a region some cells run above normal and others below, and a regional mean washes towards zero, while dry-season rain comes from large coherent systems that push a whole region positive at once. A half-degree cell is about fifty-five kilometres across and averages a flood zone with unaffected neighbours, and station-interpolated precipitation products are known to smooth extremes (Dell et al., 2014). Monthly aggregation dilutes a ten-day event across thirty-one days. Consistent with all three, only 12 percent of the fifty wettest months on record fall in the monsoon against a 33 percent null expectation, so the index's wet extremes are actively skewed away from the season when India floods. A paper about flooding specifically would need an event database or satellite inundation mapping

in place of a monthly water-balance index.

# E Additional Discussion

This appendix holds the discussion that Sections 4 and 5 refer to and that the main text does not need in full.

*The Retired Placebo in the Wet Tail.* The wet-tail income estimate for retired households, −0.0148 with a standard error of 0.0076, is out of place, since a retired household has no wage income for a stopped work site to remove. Three explanations are possible, and the data can test only the second, and only partly. Pension and transfer disbursement may itself be disrupted by heavy rain. Household composition may change, with working-age members arriving or leaving, so that a household classified as retired at baseline no longer is. Or the wet tail may be correlated with district-specific seasonality that household and state-by-block fixed effects do not absorb, in which case the specification is at fault and not the mechanism. The second has some support, since re-estimating on retired households whose member count did not change between waves shrinks the estimate by about two fifths, to −0.0091 with a standard error of 0.0080 (Appendix Table A12), without removing it.

*Departures from the Ordering.* Cereals rises slightly in rupees after a dry shock in the all-household estimate, which is consistent with substitution towards staples and equally with the higher cereal prices a dry month brings, and falls by nearly one percent more where exposure was higher, imprecisely. Intoxicants, often assumed to be among the first things cut, fall only 3.6 percent, less than sweets, where the stress consumption literature would predict otherwise, and one reading is that tobacco and alcohol are far less deferrable for the households that buy them than the discretionary label suggests, which we cannot test directly. The wet tail produces a similar ordering with two departures. Transport falls considerably more, 4.9 percent against 2.8, which fits stopped work and flooded roads, and vegetables rise instead of falling, which in rupees may be price.

*Borrowing from Shops.* Borrowing from shops is the largest single source of household credit in the survey, reported by 21.4 percent of households, ahead of banks at 18.4 percent, and the survey records it only when a payment to a shop is overdue by two months or more. If displaced borrowers moved to whatever informal source is nearest, overdue shop accounts should be among the first margins to rise when a formal lender withdraws. The row does not rise where exposure was higher, −0.07 against 0.44 where it was lower, a difference of −0.51 percentage points, though the standard error of 0.76 leaves the sign open. Two interpretations are possible, and we cannot distinguish between them. Shop credit may be trade credit from a merchant whose own working capital was financed by the same non-bank lenders, in which case the contraction hit the supplier along with the household and the informal channel is blocked at its source. Or an overdue account may be tied to purchase volume, so that a household cutting its discretionary

basket mechanically runs up less on account. The first would extend the credit-supply mechanism in an interesting direction, and the second would mean the margin is not a coping instrument at all.

*Why Displaced Households Did Not Move to Banks.* Three readings are possible, of which the first is that the households a non-bank served may not be bankable at all, in which case the two lender types serve different borrowers and the share shift in Figure 1 reflects banks lending more to their own customers rather than absorbing anyone else's. The adjustment may take longer than the four blocks the panel allows, since a bank relationship is built on a repayment record and a household displaced in late 2018 would still have been building one in 2019. Or displaced households may have moved through informal sources on their way to a bank, of which the transition table gives no sign, at a precision that admits a rise of a quarter of the base. The first reading would change how the retail credit market ought to be described, with banks and non-banks as lenders to different populations and not competitors for the same one. The bureau offers one clue, since the loans that did grow in exposed districts were the small loans of banks and financial-technology lenders (Table 1), and the survey records fintech lenders under the same non-bank label as the lenders that withdrew, so some of the displaced households may reappear in the survey's non-bank row with a different lender and a smaller loan, which would make the fall in that row an understatement of the withdrawal.

*Credit against Public Transfers.* We estimate, on the exposed occupations,

$$Y_{hdt} = \beta_s \, S_{dt} + \kappa \left( S_{dt} \times \text{Urban}_d \right) + \sum_j \theta^j \left( S_{dt} \times B_d^j \right) + \alpha_h + \tau_{st} + \varepsilon_{hdt} \tag{11}$$

where $B_d^j$ indexes three district-level buffers frozen before the panel and standardised to unit variance so that the coefficients are on one scale, with their levels absorbed by the household effects. Formal credit access is the district's bureau originations per person over 2015 and 2016, and the two transfer buffers are the share of the district's surveyed households that received a government transfer, and a private transfer, over the first two waves of the panel, interviewed in 2014, computed leaving each household's own receipt out. A buffer measured in the block before the shock is itself an outcome of past weather, since Table 2 shows originations answering a dry month, which is why the frozen measures are the ones reported. The specification carries the district's urban share interacted with the shock, $\kappa$, and a weather response that varies by state, $\beta_s$, so that the credit buffer is not urbanisation by another name. The coefficients $\theta^j$ are the difference in the discretionary-share cut per standard deviation of each buffer, so a positive value means the cut is smaller where the buffer is larger, and the table describes heterogeneity in the weather response and does not run a race, since the two transfer buffers are small-sample shares from the survey's own districts, a hundred of which hold three or fewer sampling units, and are measured with more error than the bureau's near-universe. Table A15 reports the result, and the cut is 0.31 percentage points smaller per standard deviation of formal credit when the

three buffers enter together, and 0.36 on its own, 0.12 points smaller per standard deviation of government transfer coverage and 0.18 per standard deviation of private transfer coverage, and neither transfer estimate is distinguishable from zero or from the credit estimate. Panel B asks the same question inside the design, by estimating $\theta$ separately for households that received a government transfer in the preceding wave and for households that did not. The credit effect is smaller for recipients, $-0.66$ against $-0.97$, and the difference, 0.31 with a standard error of 0.74, is not established. One caveat matters for the comparison, which is that the survey's transfer variable excludes wages from the rural employment guarantee, which are recorded as labour income. Public works scale up exactly when dry seasons hit and are known to move private wages (Imbert and Papp, 2015), so this is a race between credit and cash transfers, not between credit and the whole public safety net, and if public works availability is correlated with non-bank reliance across districts the credit estimates partly absorb the employment guarantee.

**Table A15:** Credit Access and Public Transfers Compared

| Buffer | Estimate (pp) | Std. error |
|---|---|---|
| *Panel A. Difference in the dry-shock cut to the discretionary share per standard deviation of each baseline buffer, equation* (11) | | |
| Formal credit access (district originations per person, 2015 and 2016) | 0.31** | (0.14) |
| Government transfers (district share of households receiving, waves 1 and 2 of 2014, leave-one-out) | 0.12 | (0.11) |
| Private transfers (district share of households receiving, waves 1 and 2 of 2014, leave-one-out) | 0.18 | (0.13) |
| Formal credit access, entered on its own | 0.36*** | (0.13) |
| *Panel B. $\theta$ from equation* (4) *by transfer receipt in the preceding wave, exposed occupations* | | |
| Households receiving a government transfer | -0.66 | (0.55) |
| Households not receiving one | -0.97** | (0.49) |
| Difference | 0.31 | (0.73) |
| Household-blocks | 863,296 (exposed occupations) | |

Panel A reports estimates of $\theta^j$ from equation (11) on the exposed occupations, the difference in the dry-shock cut to the discretionary budget share per standard deviation of each buffer, in percentage points, over the full estimation sample. The first three rows enter the three buffers together and the fourth enters formal credit access on its own. Each buffer is a district-level measure frozen before the panel and standardised, with the transfer shares computed leaving each household's own receipt out. Panel B reports $\theta$ from equation (4) on the exposed occupations that did and did not receive a government transfer in the preceding wave. Standard errors clustered by district in parentheses. *, ** and *** denote significance at the 10, 5 and 1 percent levels.

*Benchmarks from the Literature.* Appendix Table A13 collects the closest published estimates, beginning with Jack and Suri (2014), who find that a negative shock lowers the consumption of Kenyan households without mobile money by about 7 percent and leaves the consumption of users unaffected, Lane (2024) finds that Bangladeshi farmers offered a guaranteed emergency loan are less severely affected when a flood arrives, and Suri et al. (2021) find that Kenyan households with access to a digital loan are less likely to forgo expenses after a negative shock. In each case the outcome is the level of consumption or an expense forgone and the treatment is the addition of access, where ours is the composition of the basket and the removal of access. Breza and Kinnan (2021) find that a one standard deviation rise in district exposure to the Andhra

Pradesh microfinance crisis lowered the casual daily wage by 3.4 rupees, weekly household labour earnings by 45 rupees, and monthly household consumption by 151 rupees, with the effects confined to rural markets and small towns. Cramer et al. (2025) find that the response of non-bank lending to a weather shock falls by 1.7 log points after the IL&FS default, with the fall concentrated in collateralised products, while the response of financial-technology lenders is unchanged. Malaeb et al. (2026) supply the closest household comparison on the composition margin, a flat monthly total with non-food spending cut, on a different rural panel and a different rainfall construction, and we read the agreement as evidence that the pattern is not an artefact of one instrument.