Community-Based Dissemination Transforms Flood Forecasts into Protection: Experimental Evidence from India*

Maulik Jagnani† Tufts

Rohini Pande‡ Yale

This Version: September 2026

Abstract

Flood early warning systems promise low-cost protection against the most widespread climate hazard, but warnings work only if they reach households in time to act. Over four monsoon seasons, we run a randomized experiment across 319 communities in Bihar, India, in which trained community agents relay existing flood forecasts to households. Treatment raises the share of households receiving an alert by 9 percentage points, more than triples the number of alerts received, and increases advance protective action. Welfare gains concentrate in severely flooded villages, where treatment raises a post-flood welfare index by 0.3 standard deviations, offsetting most of the loss severe floods impose on control households. There, treated households experience less home flooding and food insecurity, and cut distress borrowing by ₹12,600, roughly a third of average seasonal borrowing. At ₹23 per household-season, averted borrowing alone implies a benefit-cost ratio of 29 to 1. A forecast protects only when it is delivered.

JEL Classification: C93, D83, H84, O12, O33, Q54

Keywords: early warning systems; last-mile delivery; climate adaptation; natural disasters; floods; information frictions; trust in technology; India

*We thank the team at Inclusion Economics: Jenna Allard, Akash Bhatt, Urmi Bhattacharya, Sathia Chakrapani, Harsh Jaiswal, and Abishek Sharma for outstanding project management; Annanya Mahajan, Advait Moharir, Priti Rajagopalan, Hamid Hussain Varikodan, Mithlesh Verma, Rohit Walthare, Roshni Yadav, and especially Surya Ravindran and Rajesh Sarma, for excellent research assistance in the field; Vineet Gupta, Miray Omurtak, Samira Salwan, George Taylor, Satish Wasti, and Aaron Wolf for excellent research assistance; and Shouryamaan Banerjee, Kairav Bhalla, John Demianczuk, Brooke Istvan, Sonia Jin, Ishan Kalia, Anuj Khanna, Rada Pavlova, Kunal Peety, Maike Pfeiffer, Brian Wong, and Adriana Zhou for valuable research support. We are deeply grateful to our NGO partner Yuganter, especially Sanjay Pandey and Amar Sharma, in helping us implement the intervention. We thank Micah Berman, Alex Diaz, Oren Gilon, Justin Ginnetti, Brigitte Gosselink, Asher Metzger, Ricki Meyer, Sella Nevo, Ofir Reich, Emily Reinstein, Juliet Rothenberg, and especially Moriah Royz at Google.org for project support. We are extremely grateful to Grey Nearing at Google.org for detailed feedback on a preliminary draft of the paper. We gratefully acknowledge financial support from the King Climate Action Initiative (K-CAI) at J-PAL, Google.org, the National Science Foundation, the Agency Fund, Wellspring Philanthropic Fund, and GDI Solutions.

†The Fletcher School and Department of Economics, Tufts University, Medford, MA 02155. Email: maulik.jagnani@tufts.edu

‡Department of Economics and Economic Growth Center, Yale University, New Haven, CT 06511. Email: rohini.pande@yale.edu

1 Introduction

Floods are the most widespread and costliest of climate hazards: roughly 1.8 billion people face significant flood risk, nine in ten of them in low- and middle-income countries, and riverine flooding alone causes average annual losses of roughly $388 billion.[1] Against this risk, early warning systems (EWS) are a widely promoted adaptation investment. Projected benefits exceed costs many times over, disaster mortality is nearly six times higher where coverage is limited than where it is comprehensive, and the UN's Early Warnings for All initiative targets covering every person on Earth by 2027.[2] The case for EWS strengthens as forecasting advances rapidly, through new country-level systems and new forecasting technologies deployed at global scale (Nearing et al., 2024). But a forecast protects households only if the warning reach them in time to act, a margin typically assumed away when projecting the returns to EWS (e.g., Pappenberger et al., 2015; Teisberg and Weiher, 2009).

Locating this margin requires unpacking the warning chain. A stylized chain has five links, forecast $\rightarrow$ guidance $\rightarrow$ alert $\rightarrow$ dissemination $\rightarrow$ action (adapted from Reichstein et al., 2025), running from hazard prediction through warning delivery to protective action by governments, response agencies, or residents. A sequential chain is only as effective as its weakest link, and in much of the world the weakest link is dissemination. The number of countries reporting a multi-hazard EWS has more than doubled in the past decade (United Nations Office for Disaster Risk Reduction and World Meteorological Organization, 2024), but a system on paper is not a warning received: in a recent survey across 142 countries, 30 percent of people who experienced a disaster received no warning at all, with rural and less educated residents warned less often (Lloyd's Register Foundation, 2024), and in South Asia riverine flood warnings reach fewer than half of households in exposed areas (Lang et al., 2025). Reviewing this century's deadliest and costliest weather disasters, Coughlan de Perez et al. (2022) conclude that the missing protection was not a lack of forecasts but a lack of communication and response capability, so that better forecast accuracy will yield little until these downstream gaps close.

Why does the last mile fail even where forecasts are accurate? Our setting illustrates the mechanics. The two systems covering Bihar predict floods accurately two to five days ahead, but neither delivers to households. India's Central Water Commission (CWC) broadcasts through mass media and websites and relays alerts to local officials, not to residents. Google pushes notifications directly, but only to Android phones with location services enabled, a setting most smartphone owners in Bihar do not use. The result is that forecasts covered over 90 percent of our control households, yet in a typical season only 35 percent received any warning, and 44

[1] Rentschler, Salhab and Jafino (2022); United Nations Office for Disaster Risk Reduction (2025)

[2] Global Commission on Adaptation (2019); United Nations Office for Disaster Risk Reduction and World Meteorological Organization (2024); World Meteorological Organization (2022).

percent experienced a flood with no alert. Unreliable delivery then erodes trust: 87 percent of households that experienced a missed flood or false alarm say it reduced their willingness to act on future warnings.

In this paper, we quantify the value of the dissemination link: does effectively delivering an existing flood forecast induce protective action and reduce the ex post costs of floods? In partnership with a Bihar-based NGO, we randomly assigned 160 of 319 gram panchayats (GPs), the lowest tier of local government and home to 3.6 million people, to a multi-channel delivery program. The GPs span twelve districts of North Bihar, the floodplain north of the Ganges in India's poorest state (NITI Aayog, 2023), where three-quarters of the population lives under recurring flood threat (Water Resources Department, Government of Bihar, n.d.). In treated GPs, trained local agents monitored official forecasts, from India's Central Water Commission (CWC) from 2022 and Google Flood Hub from 2023, and relayed warnings through loudspeakers, flags, SMS, and WhatsApp. Because the program produced no new forecasts, the experiment isolates the dissemination link.

The experiment spanned four flood seasons (2022 to 2025). We track a panel of 5,582 households, roughly 21,000 household-season observations, paired with a satellite-based measure of flood severity for each village and season and an incentivized elicitation of willingness to pay (WTP) for the service. This design lets us estimate causal effects on (i) alert receipt and preparatory action, (ii) post-flood welfare, including asset losses, food security, and distress borrowing, and (iii) beliefs about warnings and demand for the service. A warning is a state-contingent good: it is valuable only when the hazard materializes, and more so the more severe the hazard. Pooled treatment effects therefore understate its value in the states it exists to insure. We estimate effects on post-flood welfare as a function of realized flood severity, which varies across villages and seasons.

First, the program delivered warnings, and household acted on them. Treatment raised the probability of receiving any flood alert by 9.3 percentage points on a control base of 35 percent and more than tripled the number of alerts received, with gains concentrated in the agents' direct-to-household channels, SMS and WhatsApp. Quality improved alongside quantity: alerts in treated GPs arrived more often before floodwaters reached the home and were more often rated completely trustworthy, very useful, and very accurate, and treated households more often passed flood information on to community members. Measured against self-reported flooding, received warnings became more accurate on the margin that matters: dangerous misses, floods with no warning, fell by 6.1 percentage points from a 43.9 percent control base, at the cost of a 3.5 percentage point rise in false alarms (from a 5.2 percent base), a net reduction in forecast errors of 2.6 percentage points. Households converted the added lead time into preparation: an index of advance protective actions (including stocking food, sandbagging, and digging drainage chan-

nels) rose by 0.052 standard deviations, or 0.57 among households induced to receive an alert, while last-minute responses were unchanged.

Second, preparation paid off when and only when the floods arrived. The treatment effect on post-flood welfare index tracks the loss it insures against: each additional day of flooding lowers a control household's welfare index by 0.018 standard deviations, treatment offsets 0.015 of that loss per day, and full offset cannot be rejected ($p = 0.31$). Under severe flooding, the treatment effect of 0.30 standard deviations recovers most of the -0.36 penalty; with no flooding the effect is zero. In severely flooded villages, gains are visible in the components: treated households were 13 percentage points less likely to have floodwater enter the home and avoided four-fifths of the house damage severe floods cause, were 7 percentage points less likely to be food insecure, and borrowed ₹12,632 less, roughly 40 percent of flood-season income. None of this drew on extra public support: treated households received no more cash transfers or rations, and relied less on emergency community kitchens.

Third, effective warning delivery shifted beliefs and generated demand for the service. Control households in villages that were flooded more often were more likely to be worried about unexpected future flooding, but treatment offsets this gradient. In severely flooded villages the season's experience improved views of science and technology. Updated beliefs raised willingness to pay: in an incentivized Becker-DeGroot-Marschak elicitation, treatment increased WTP for delivery by 5 percent, and by 37 percent among households induced to receive an alert. Even control households' average WTP of ₹34 per season exceeds the ₹23 delivery cost, and counting only averted distress borrowing, benefits exceed its costs 29 to 1.

Contributions. First, we provide what are, to our knowledge, the only experimental estimates of the welfare effects of delivering disaster early warnings. Existing EWS benefit figures rest on simulation and avoided-damage accounting (Pappenberger et al., 2015; Teisberg and Weiher, 2009), and evidence on forecast value in lower-income countries comes from agriculture, where forecasts inform seasonal planning rather than protection against an imminent hazard: long-range monsoon forecasts shift planting decisions in both observational (Rosenzweig and Udry, 2013) and experimental (Burlig et al., 2024) settings, while SMS weather and price advisories have more modest effects (Camacho and Conover, 2019; Fafchamps and Minten, 2012; Yegbemey, Bensch and Vance, 2023).[3] Interacting treatment with an objective severity measure also lets us trace the state-contingent payoff of warnings: zero in village-years without floods, near-complete offset of losses under severe flooding. A corollary is that pooled treatment effects are a misleading

[3]In rich countries, where mass media saturate the population with forecasts, quasi-experimental work values improvements in forecast accuracy (Molina and Rudik, 2024; Shrader, Bakkensen and Lemoine, 2023). Our setting inverts this: prediction is already accurate, delivery is the margin that binds, and we vary delivery experimentally while holding forecast production fixed.

summary of warning-system value. Our pooled estimate on the welfare index, 0.034 SD (s.e. 0.028), is indistinguishable from zero. It averages a large effect in severely flooded village-years with smaller effects in mildly flooded village-years and zero effects everywhere else, and, like any single-window estimate, it is conditional on the window's draw of shocks (Rosenzweig and Udry, 2020). Our window was mild: the realized incidence of severe flooding was 5.4 percent, against 15.7 percent over 2017 to 2025. The severity interaction separates the effect in each state from how often that state occurred, which is what allows reweighting to the longer-run distribution. Our cost-effectiveness calculation does not reweight; it uses the realized 2022 to 2025 frequency, which is one of the margins on which it is conservative.

Second, we engage with the literature on the last mile of technology in development. Its premise is that money and technology deliver returns only through the human systems that carry them to households, what Page and Pande (2018) call invisible infrastructure, and the returns to promising technologies have hinged on that final link in health, agricultural extension, payments and education (Björkman Nyqvist et al., 2019; Cristia et al., 2017; Dodge et al., 2025; Fabregas, Kremer and Schilbach, 2019; Hanna, Duflo and Greenstone, 2016; Muralidharan, Niehaus and Sukhtankar, 2016). Artificial intelligence sharpens this problem: as prediction becomes cheap, value shifts to its complements (Agrawal, Gans and Goldfarb, 2018). Flood forecasting is a case in point. AI-based forecasts are free and public in more than 150 countries (Google Research, 2026), yet in our study area the direct digital channel that carries them reached fewer than half of households even when a notification was issued, and 75% of control households received no notification despite floodwater entering their house. We provide, to our knowledge, the first experimental welfare estimates of the human complement to AI flood forecasting at scale: an agent layer costing ₹23 per household-season that converts existing forecasts into protection. Agents relayed official CWC and Google alerts without modification, adding only reach and, possibly, the credibility of a familiar sender; our design prices that bundle rather than apportioning value between the two. Consistent with experience driving demand for new technologies (Dupas, 2014), households revised their beliefs about warning reliability where floods verified the forecasts, and treated households were willing to pay more for the delivery service.

Third, we contribute to the evidence on ex ante protection against climate shocks, from forecast-based anticipatory action to broader social protection. Anticipatory cash protects households: transfers arriving within days of the flood peak improve food security (Christian et al., 2025; Pople et al., 2026), reduce asset losses (Pople et al., 2026), and lower post-flood borrowing (Christian et al., 2025), and even a pre-committed credit line, drawn only if a flood occurs, raises ex ante investment and post-flood consumption (Lane, 2024). At larger scale, anti-poverty programs themselves function as ex ante adaptation: the assets and occupational diversification induced by BRAC's graduation program largely offset the consumption losses that unpredictable floods and

droughts impose on ultra-poor households (Balboni et al., 2026). In the framework of Carleton et al. (2024), these interventions relax resource constraints, through liquidity timed to the shock or assets accumulated before it, while holding information fixed or bundling it with transfers. We instead vary the information input on its own. Warning delivery costs ₹23 per household-season, under one percent of the anticipatory transfers evaluated in this literature and orders of magnitude below graduation-program outlays, yet it moves the same family of outcomes when severe floods strike: food security, asset losses, and distress borrowing. Households financed their own preparation, at outlays of one and a half to four times those transfers,[4] so a transfer is not a precondition for anticipatory protection in this setting. Our relief results add a financing margin: treated households drew less on community kitchens during floods, with no detectable change in government cash relief, evidence that information delivered ex ante can substitute for a component of relief ex post.

2 Background and Research Design

2.1 Floods in North Bihar

Bihar is India's poorest and most flood-prone state, accounting for more than one-fifth of its flood affected population (District Nalanda, Government of Bihar, n.d.; National Remote Sensing Centre, 2020). Himalayan rivers drain from Nepal's steep catchments onto a flat, heavily silted alluvial plain, spreading monsoon floods laterally across vast areas: 73% of the state's area is flood-affected, triple since the 1950s, and roughly 76% of North Bihar's population lives under the recurring threat of severe flooding (Bandyopadhyay, 2017; Water Resources Department, Government of Bihar, n.d.). Our satellite mapping of inundation over 1984–2021 confirms that flooding has intensified on both the extensive and intensive margins (Figures A1–A2).

The exposed households are poorly placed to absorb these losses. In a survey in seven districts in north Bihar, 87% lived in kucha dwellings made up of mud and other non-durable materials, and the 2024 floods forced widespread distress coping: nine in ten households cut meals, two-thirds borrowed from moneylenders at high interest (Prasad, Verma and Patil, 2025). The same assessment finds that flood warnings "rarely translated to actionable information that flood vulnerable communities could act upon:" the gap between forecast production and household protection that this study addresses.

[4] Among preparing households in flooded areas in 2025, mean preparation expenditure was roughly ₹14,000. Transfers were BDT 4,500 (about ₹4,000) in Pople et al. (2026), and BDT 5,000 and NPR 15,000 (about ₹3,500 and ₹9,400) in Christian et al. (2025).

2.2 Flood Early Warning Systems in Bihar

India operates among the region's most developed forecasting infrastructure, and our study communities are covered by two systems. India's Central Water Commission (CWC) runs a national network of 338 flood-forecasting stations, issuing over 10,000 flood forecasts per year from machine-learning-based basin-wise rainfall-runoff models with two-to-five-day lead times (Press Information Bureau, 2022).[5] Its warnings are disseminated through web portals, mass media, and push alerts to disaster-management authorities, reservoir and river-basin operators, and registered stakeholders, not to households.[6]

Google's Flood EWS piloted in India in 2018. It feeds CWC gauge readings into a hydraulic model built on high-resolution digital elevation models, translating river conditions into spatial maps of areas likely to flood.[7] Forecasts are issued two to five days ahead, and Android smartphones with location services enabled in likely affected areas are notified (Figure A3(a)).

Both systems classify flood situations into three levels (normal, above normal, severe) based on river water levels relative to pre-defined warning and danger levels (Central Water Commission, 2020). A typical event begins with a warning (yellow) alert, escalates to severe (red), and reverts to yellow as waters recede.

2.3 The Intervention: Community-Based Alert Dissemination

Our intervention concerns only last-mile delivery; the forecasts themselves are unchanged. In each of the four monsoon seasons (2022-2025), in partnership with the local NGO Yuganter, we recruited and trained community volunteers (agents) in treatment GPs to disseminate the alerts

[5]The models draw rainfall inputs, from the Global Forecast System (GFS) and Weather Research and Forecasting (WRF) models, supplied by the India Meteorological Department. Underpinning the forecasts is an extensive data-collection layer: more than 1,700 hydro-meteorological sites across 20 river basins, transmitting observations in near-real time via wireless, telephone, mobile, and satellite-based telemetry, including VSAT ground stations (Press Information Bureau, 2022). The network spans 25 states and union territories and has expanded steadily over time; 338 is the figure for recent flood seasons (Press Information Bureau, 2023).

[6]The web portal carries both a seven-day advance advisory forecast for major river basins and shorter-term current forecasts. Mass media outreach runs through print, social, and electronic media, including a dedicated flood-updates channel. These forecasts could also reach the public through the Sachet and FloodWatch India apps (Central Water Commission, 2024; National Disaster Management Authority and Centre for Development of Telematics, 2024).

[7]Where the hydraulic model's accuracy is compromised (e.g., through digital elevation model inaccuracies, embankment breaches, or unexpected water sources), Google employs a second inundation model trained on historical data from the European Space Agency's Sentinel-1 satellite constellation, which carries C-band Synthetic-Aperture Radar (SAR) instruments. SAR imagery detects inundation under all weather conditions, including through cloud cover. By correlating historical water-level measurements with past inundations captured by SAR, Google distinguishes errors caused by genuine changes in ground conditions from those resulting from model inaccuracies, enabling more reliable predictions (Google Research, 2019). We construct our own satellite-based flood-exposure measure from Sentinel-1 SAR imagery on a similar principle; see Section 4.2.

from two systems: CWCs throughout and Google's from 2023.[8] Agents, three per treatment GP in 2022, two thereafter, served from June to October, coinciding with Bihar's monsoon, for a monthly stipend of ₹2,000, and were trained each May to receive, interpret, and relay alerts. Whenever a water station linked to any village in their GP triggered a warning (yellow) or severe (red) alert, agents disseminated it through traditional channels, such as loudspeakers and flags, and direct-to-household channels: SMS and WhatsApp groups established for each treatment GP (Figure A3(b)). From 2023 onward agents also relayed Google's flood maps, attached to each alert, delineating areas likely to flood and those expected to remain dry (Figure A3(a)). Agents logged daily whether they had received alerts and what actions they took. No agents were recruited in control GPs.

2.4 Experimental Sample and Randomization

The evaluation comprises 319 GPs, home to more than 3.6 million people, across 12 flood-prone districts of Northern Bihar (Figure 1(a)).[9] GPs were selected, jointly with our NGO partner, as flood-prone and covered by CWC forecasting stations. To ensure treatment and control communities face comparable flood risk, Google's engineering team mapped every village to its relevant river gauge stations using each station's Area of Interest: the set of villages whose flood conditions that station's readings inform (Figure 1(b)). We randomly assigned 160 GPs to treatment and 159 to control, stratified by district. The two arms are well balanced on household characteristics and realized flood exposure (Section **??**). Because both arms face identical forecasts and alerts, the experiment isolates the marginal value of community-based last-mile dissemination, not the value of the forecasting systems themselves.

3 Conceptual Framework

Earlier we established that treatment varies only delivery. This section derives what improved delivery should change and why.

3.1 Information, Constraints, and Protective Action

Environment. A household enters the monsoon with flood-exposed assets: standing crops, stored food, livestock, the house, documents and valuables, and its members' health. A flood

[8] Google alerts were unavailable in 2022.

[9] Gram panchayats (GPs) are the lowest tier of administrative division in rural India; each typically oversees two to five villages.

of severity $s \geq 0$ is drawn from a distribution with mode $s = 0$; as in our data, s counts the days village inundation exceeds a threshold. Unprotected damages $D(s)$ rise steeply with severity.

Actions and lead time. Protective actions divide by the lead time they require. *Advance* actions work only if taken days before onset: harvesting early, moving crops to higher ground, digging drainage, sandbagging, vaccinating livestock, and stocking food, water, and medicine. They avert damages $\beta(s)\, x$, where x is the intensity of advance action and $\beta'(s) > 0$: the same drainage canal or early harvest saves more the longer water stands. *Last-minute* actions, such as moving people and documents, need no forecast; flood onset is self-revealing.

Information and trust. Forecasts are generated for nearly all sample villages. A forecast enters a household's information set only if it is *delivered*: received before onset, understood, and believed. Let ρ be the probability that a household receives a warning ahead of a flood, and τ the credibility it attaches to warnings, shaped by past false alarms and misses. Given its signal, the household forms a posterior $\pi(\rho, \tau)$ that a damaging flood is imminent and takes advance action if perceived expected averted damages exceed costs:

$$\pi \cdot \mathbb{E}[\beta(s) \mid \text{flood}] \geq c,$$

where c is the per-unit cost of advance action.

Two features of the cost structure matter. First, in bad states the cost of advance action is small relative to damage it averts: $c \ll D(s)$ for high s: a warned household with any means to act will act. Information, not liquidity, is the first-order constraint, though liquidity may still limit how much a household prepares. Second, the two forecast errors carry different prices: an acted-upon false alarm costs c. An unwarned severe flood costs $D(s)$. A delivery technology that converts dangerous misses into a few additional false alarms therefore raises expected welfare; the trade that public amplification of every issued alert makes.

What treatment changes. Community agents amplify existing warnings: they raise ρ, deliver alerts earlier, and, through repeated and publicly verifiable delivery, raise τ. The forecast itself is unchanged. The intent-to-treat effect is the product of the share of households newly reached and the response of a household to a warning, and so scales with the first stage. This motivates instrumental-variables estimates of alert receipt for outcomes that are a household's own decisions, advance action and willingness to pay. We explain below why the same estimator is not appropriate for welfare.

Prediction 1 (delivery margins). *Treatment increases the receipt, frequency, and timeliness of alerts and shifts delivery toward agent channels; it sharply reduces missed alerts while modestly raising false alarms; and it raises trust in warnings.*

3.2 Why Effects May Concentrate in Severe Floods

State-contingent payoffs. Advance action is an insurance premium which is paid whenever an alert is acted upon, but redeemed only if the flood materializes. Conditional on realized severity, the treatment effect on welfare is

$$\Delta W(s) \;=\; \big[\beta(s) - c\big]\,\Delta x,$$

where Δx is the treatment-induced increase in advance action and $\beta(s)$, weakly increasing in s with $\beta(0) = 0$, is the damage averted per unit of action. Four implications follow. In unflooded villages $\Delta W(0) = -c\,\Delta x \leq 0$: the premium paid on protection that was not needed is small if the cost is small or recoverable; a positive effect there would point to a mechanism other than the one modeled. The effect increases with s, a dose-response in realized severity. Gains load on the domains advance action protects (assets, harvests, food stocks) and the ex post coping those protections substitute for; entitlement transfers, which depend on flood incidence and administrative targeting rather than on the household's own preparation, should be unaffected. And, the pooled treatment effect is the incidence-weighted average, $\mathbb{E}[\Delta W] = \sum_k p_k\,\Delta W_k$ across severity bins.

Severe floods. Physical damage is convex in duration: losses spread from housing to standing crops, stored food, and land as water stands longer (Gupta et al., 2026; Li, Rosa and Gorelick, 2025; Merz et al., 2010; Prasad, Verma and Patil, 2025). If advance action averts a share of this damage that does not shrink as floods worsen, $\beta(s)$, is convex too, and the return to a delivered warning rises faster than severity. For low-severity events, priors and local observables such as rainfall and river levels already support adequate preparation such that a delivered forecast adds little.

Coping costs and the extent of protection. These forces raise the treatment effect in severe floods, but not necessarily the *share* of the loss offset, since the loss itself increases. Coping costs inform the impact on the share. Let $L = D(s) - \beta(s)\,x$ be the loss remaining after protection, and $K(L)$ its coping cost. Households hold a buffer: savings, stored food, help from relatives. Within the buffer, coping is cheap; beyond it they cut meals and borrow at high interest, so K is near zero for small L and rises steeply past the buffer. The treatment effect becomes

$$\Delta W(s) \;=\; \big[\beta(s) - c\big]\,\Delta x \;+\; \big[K(L_C(s)) - K(L_T(s))\big],$$

where L_C and L_T are the remaining losses in control and treated households. In a short flood, both losses stay within the buffer and the coping term is zero. In a severe flood, the control household's loss passes the buffer while partial protection keeps the treated household's inside

it: food insecurity and distress borrowing disappear even though asset and harvest losses are only partly averted. A sharp buffer makes the logic cleanest, but the same ordering holds whenever K is convex.

Prediction 3 (state-contingent welfare). *Treatment effects on welfare are approximately zero in unflooded villages, increase with realized severity, and are concentrated in assets, harvests, food security, and reduced distress borrowing. The share of the flood loss that treatment offsets also rises with severity, with the low-to-severe difference concentrated in food security and coping rather than physical losses.*

Prediction 4 (relief margins). *Treatment leaves entitlement transfers (cash, dry rations) unchanged while reducing reliance on emergency communal provision (community kitchens) in flooded villages.*

Learning and demand. Delivered warnings also change beliefs. Households learn the service's value experientially: the informative event is a warning validated by a realized flood and by visibly averted damages, their own and their neighbors'. Updating should therefore concentrate where welfare gains do, in treated villages that flooded severely, appearing as less worry about unwarned future flooding, fewer reported misses, and greater trust, and may extend to generalized beliefs about science and technology where technology demonstrably paid off. Treatment should raise willingness to pay for the flood-alert delivery but not for unrelated goods, distinguishing learning from income effects or experimenter demand.

Prediction 5 (learning and demand). *Trust, reduced flood worry, and beliefs about science and technology improve most where treatment coincided with severe flooding; willingness to pay rises for flood-alert delivery but not for placebo goods.*

4 Data and Empirical Strategy

4.1 Timeline, Surveys, and Sample

The experiment spans four consecutive monsoon flood seasons, 2022 through 2025, each constituting one round. Data collection is organized around two household surveys per round (Figure 2): a midline phone survey administered during the flood season itself, to improve recall, captures alert receipt, trust in warnings, and the timing of household actions; an in-person endline survey after the flood season measures realized flood exposure and downstream welfare outcomes: asset damage, health, harvests, food security, and coping. The midline thus records alert receipt and ex-ante responses as they happen; the endline records ex-post outcomes.[10]

[10]Households that missed the midline answered the relevant questions retrospectively at the subsequent in-person endline. The 2022 midline, conducted before the panel was assembled, drew on a separate sample of residents referred by panchayat leaders; we include it in the analysis of ex-ante outcomes. From 2023 onward the midline was administered to the tracked panel. In 2024 and 2025 we also included a GP-level survey. Appendix 11.1.1 provides

The tracked panel comprises 5,582 households across 591 villages in the 319 study GPs, enrolled at the first round's endline (December 2022 to March 2023).[11] We randomly selected two villages per GP and, within each, enrolled ten households by a random walk from the village center, approaching every third household, subject to two quotas: at least eight of the ten households owned a smartphone, and Scheduled Caste and Scheduled Tribe households were represented in proportion to their village population share. Appendix 11.1.2 details the protocol, including the replacement rule when a quota would otherwise have gone unmet. Oversampling smartphone-owning households gives the status quo's direct alert channels their best chance in both arms: these are the households most reachable by Google's Android-based alerts, so treatment effects measure the value of community-based dissemination over and above the direct channel.

The sample is poor, agricultural, and exposed (Table 1; we describe control means). The average household has 5.5 members, earns about ₹31,000 over the flood season, owns just under one acre of land, and lives 2.8 km from the nearest river; 68% report agriculture as an occupation, 41% live in permanent dwellings, and four-fifths of households belong to historically disadvantaged groups (21% SC/ST, 59% OBC/EBC). By design, smartphone ownership is near-universal at 91.5%.

4.2 Measuring Flood Exposure and Realized Severity

Our analysis requires a measure of realized flood exposure that is granular, available for every village in every season, and, critically, independent of treatment itself. Satellite remote sensing satisfies all three: it measures inundation identically, and mechanically, in treatment and control villages (Patel, 2024, 2025). We construct a daily, village-level inundation share, the fraction of village area under water, for every June–October monsoon season from 2017 through 2025, from Sentinel-1 Synthetic Aperture Radar (SAR), which detects standing water at night and through cloud cover, gap-filled between satellite passes by imputation.[12] Appendix 11.1.3 details the sensor, the imputation model, and its validation.

Two steps convert this series into exposure. First, because the raw satellite measure records standing water of any kind (rivers, ponds, routine monsoon waterlogging), we residualize the daily share on village-by-week-of-year and village-by-year fixed effects estimated on the full 2017–2025 panel, so the measure captures *anomalous* inundation relative to each village's own seasonal

details.

[11] Given the uncertain trajectory of the COVID pandemic, we chose not to conduct a pre-intervention baseline survey. Time-invariant characteristics collected at this first endline serve as the baseline for the balance tests in Section **??**.

[12] Because the satellites' revisit cycle yields direct observations for only about 19% of village-days, we fill the remainder with a Random Forest imputation trained on the observed days. Out of sample it achieves an R^2 of 0.93 with a mean absolute error of 1.8 percentage points of village area, and 0.92 and 3.2 percentage points in monsoon months, when inundation is most variable.

rhythm.[13] Second, we define a *flood day* as a village-day on which anomalous inundation exceeds 30 percentage points of village area, an order of magnitude above the imputation model's monsoon-month error and physically unambiguous: nearly a third of a village newly under water. The count of flood days in a village-year is our continuous severity measure S_{vt}; for binned specifications we classify village-years as *no flooding* ($S_{vt} = 0$), *low flooding* ($1 \leq S_{vt} \leq 7$), and *severe flooding* ($S_{vt} > 7$). Results are robust to alternative thresholds and to episode-based duration measures (Appendix 11.1.3).[14]

Applying this classification, 77% of village-years experience no anomalous flooding, 17% low, and 6% severe; the quarter with some flooding is where flood-conditional effects are identified. The four seasons trace out the severity distribution: 2023 was nearly floodless (13% of villages flooded, exactly one severely), 2022 moderate, 2024 substantial, and 2025 the floodiest (33% flooded, 15% severely). Two features matter for interpretation. First, severe exposure concentrates in the later rounds (129 of 147 severe village-years occur in 2024 and 2025), but because every specification includes district-by-round fixed effects, severe-flood effects are identified by comparing treated and control villages facing the same season within the same district, not by comparing flood years to mild years. Second, the panel contains substantial within-village variation: 104 of 591 villages flood severely in at least one season and none in all four, so every severely flooded community is also observed under milder realizations. The window is mild by local standards: in the five preceding monsoons, 23% of village-years were severe, nearly four times the study-window rate. Households therefore entered the experiment with recent, first-hand experience of severe flooding, and our pooled estimates are, if anything, conservative relative to the floodier seasons that Bihar's rising flood incidence makes likely. Appendix 11.1.4 reports season-by-season incidence, within-season timing, and village flood histories.

4.3 Outcome Variables and Index Construction

Outcomes are organized along the causal chain the experiment traces: whether warnings reach households, whether households act on them in advance, whether losses are averted once floods arrive, and whether the experience changes beliefs and demand. From the midline we measure delivery (whether any alert was received, how many, their timeliness, mode, and informational content); quality (crossing receipt against self-reported flooding yields a *false alarm*, alert but no flood, and a *missed flood*, flood but no alert, the dangerous error an EWS exists to prevent); and

[13] On an ordinary day about 8% of the average study village's area is under water. This is water households already anticipate, so a village whose fields routinely hold water in August is not coded as flooded.

[14] Because a single flood can register as several bursts of flood days separated by brief recessions, we also construct episode-level measures: grouping flood days separated by at most seven dry days into single events, we record each village-year's number of flood events and their average and longest durations. In practice the distinction is minor: total episode duration and the raw flood-day count correlate at 0.98 across village-years.

credibility (trust in and perceived usefulness of alerts). The midline also records protective steps against a structured list, which we partition into *advance responses*, twenty-seven preparatory steps across five domains (human health, livestock, house and environment, agriculture, and food and water), and *last-minute responses*, five actions at flood onset such as moving household members and documents. If warnings relax an information constraint, action should shift into the first category without increasing the second.

From the endline we construct five domains of realized welfare (physical health, asset damage, agricultural harvest, food security, and coping), along with receipt of government relief: dry rations, Bihar's ₹ 7,000 post-flood cash transfer, and community-kitchen meals, measured both as outcomes and to interpret welfare effects net of the government's response. The 2024 and 2025 endlines add a science-and-technology beliefs module adapted from the World Values Survey. The 2025 endline adds questions on whether in recent years households have experienced false alarms or missed alerts and how concerned they are about unexpected flooding in 2026, and an incentivized Becker–DeGroot–Marschak elicitation of willingness to pay for alert delivery, benchmarked against a non-incentivized practice round and an incentivized monsoon-onset forecast (Dizon-Ross and Jayachandran, 2022).

Our primary outcomes are summary indices, which discipline multiple-hypothesis concerns and improve power against effects spread across components (Kling, Liebman and Katz, 2007). Each component is standardized against the control group within survey round and signed so that higher values are better; the domain index is the equal-weighted average of its components, and the headline advance-response and welfare indices average their five domains, restandardized against control. A treatment effect on an index is read in control standard deviations.

4.4 Balance and Attrition

Randomization delivered balance. Across twenty-two baseline characteristics, no treatment–control difference is significant at the 5% level, magnitudes are uniformly small, and a joint test does not reject (F = 0.674; Table 2). Realized flood exposure is likewise balanced, 6.6% of control and 5.4% of treatment household-rounds severely flooded ($p = 0.34$), with mean flood-days indistinguishable ($p = 0.42$), confirming the exogeneity of severity to treatment.[15] Attrition is low and uncorrelated with treatment: we reach 88–94% of baseline households in each follow-up round, endline completion rises from 80% in 2023 to 90% thereafter, and the treatment–control difference in completion is insignificant in every round on every margin ($p \geq 0.27$; Appendix 11.2 and Table A3).

[15] In principle, treatment-induced adaptation, particularly at the community level, could marginally reduce measured standing water in treated villages. However, neither the severity bins nor total flood-days differ by treatment status.

4.5 Estimating Equations

We report three sets of estimates: pooled intention-to-treat (ITT) effects on alert receipt and ex-ante adaptation, outcomes that respond to forecasts rather than to realized flooding; treatment effects on ex-post welfare outcomes allowed to vary with realized severity, in continuous and binned specifications; and, for ex-ante adaptation and willingness to pay, instrumental-variables estimates that rescale the ITT by the program's effect on alert receipt. Throughout, treatment T_g is randomly assigned at the GP level, stratified by district, and standard errors are clustered at the GP level (319 clusters).

First stage and ex-ante outcomes. For alert delivery and ex-ante adaptation, we estimate

$$Y_{igt} \;=\; \beta\, T_g \;+\; \lambda_{s(v,t)} \;+\; \delta_{dt} \;+\; \varepsilon_{igt}, \tag{1}$$

where Y_{igt} is an outcome for household i in GP g surveyed in round t, δ_{dt} are district-by-round fixed effects aligned with the stratified assignment, and $\lambda_{s(v,t)}$ are fixed effects for realized flood days in the household's village v. Ex-ante outcomes enter uninteracted with severity: preparation is chosen under uncertainty, in response to forecasts issued before flooding is realized. The flood-days fixed effects absorb residual variation in realized conditions and improve precision; they do not model effect heterogeneity.[16]

Heterogeneous impacts by severity. Ex-post welfare effects should depend on realized flooding, and we measure severity as duration. The continuous specification interacts treatment with severity,

$$Y_{igt} \;=\; \beta_1\, T_g \;+\; \beta_2\, S_{vt} \;+\; \beta_3\, \left(T_g \times S_{vt}\right) \;+\; \delta_{dt} \;+\; \varepsilon_{igt}, \tag{2}$$

where β_2 is the loss per additional day of flooding among control households and β_3 is the dose-response slope, the change in the treatment effect per additional day. A test of $\beta_2 + \beta_3 = 0$ is therefore a test of whether treatment fully offsets the per-day loss. The binned specification relaxes linearity,

$$Y_{igt} \;=\; \sum_{k \in \{L,\, Sev\}} \gamma_k\, B^k_{vt} \;+\; \sum_{k \in \{N,\, L,\, Sev\}} \tau_k\, \left(T_g \times B^k_{vt}\right) \;+\; \delta_{dt} \;+\; \varepsilon_{igt}. \tag{3}$$

Treatment is fully interacted with the bins, so each τ_k is the treatment effect within severity bin k, and $\gamma_k + \tau_k = 0$ tests full offset within bin k. The no-flood effect τ_N serves as a falsification test for damage outcomes. We also show that heterogeneous effects on the aggregate welfare index are robust to the inclusion of village fixed effects.

[16] For first stage and ex-ante outcomes, which pool midline and endline modules, we also include a midline indicator.

Instrumenting alert receipt. The ITT averages over households whose alert receipt the program did and did not change. For ex-ante adaptation and willingness to pay, we also instrument an indicator for receiving any flood alert, A_{igt}, with assignment:

$$A_{igt} \;=\; \pi\, T_g \;+\; \lambda_{s(v,t)} \;+\; \delta_{dt} \;+\; \nu_{igt}, \tag{4}$$

with the second stage replacing T_g in equation (1) with instrumented receipt. For willingness to pay, elicited once at the 2025 endline, the endogenous variable is receipt in 2025 and the fixed effects are those of the willingness-to-pay specification. The instrument is strong: as we discuss below, assignment raises receipt by 9.3 percentage points on a control base of 35 percent, with Kleibergen–Paap F statistics between 66 and 71 across estimation samples. The exclusion restriction, that assignment affects outcomes only through the binary receipt margin, is a strong assumption in our setting: the program also made alerts earlier, richer in content, and more trusted, and may have induced village-level responses that benefit non-recipients. We therefore interpret the two-stage least squares estimates as the ITT rescaled by the compliance margin; to the extent intensive-margin improvements also raise outcomes among households who would have received an alert anyway, they overstate the effect of receipt alone, and we flag this where relevant.[17]

5 Forecasts and The Last-Mile Gap

This section documents the gap between forecast accuracy and forecast delivery, using administrative data from the two warning systems operating in Bihar and household surveys from our sample. Forecasts are accurate days in advance. Under status-quo dissemination, however, notifications reach at best about half of at-risk households, and most households in flooded communities report receiving no warning at all.

5.1 Forecast Accuracy

Figure 4 compares forecasted and realized water levels for the two forecasting systems described in Section 2, at the gauge stations serving our sample, by season and source. Each panel restricts attention to station-days on which the realized level exceeded the station-specific warning threshold, and plots binned means of the forecasted level against realized level, with the dashed line marking perfect prediction. Observations lie close to that line throughout: the correlation be-

[17] For ex-post welfare we report ITT effects only in the main text, with IV estimates in the appendix for completeness. Welfare depends on what neighbors and the community do as well as on the household's own preparation, so the exclusion restriction is unlikely to hold.

tween forecasted and realized levels is 0.96 to 0.97 for CWC (2022 to 2024), and for Google rises from 0.87 in 2023 to 0.89 in 2024 and 0.97 in 2025. (Google data are unavailable for 2022, and CWC data for 2025.) Forecast uncertainty widens at the highest water levels, where observations are sparse, and both systems there tend toward over-prediction, the conservative direction for a warning system, since a false alarm is arguably less costly than a missed flood.

Forecasts are also accurate on timing. Figure 5 traces a single station through the 2025 monsoon, the floodiest of our study seasons, plotting the daily share of its villages classified as flooded by our satellite measure alongside realized and forecasted water levels above the severe threshold.[18] The three series move together: forecasted exceedance tracks realized exceedance, village-level inundation rises and falls with the gauge, and the season's severe-alert events, from Google and from the treatment system alike, were issued before local inundation peaked. Warnings therefore arrive early enough to act on.

This accuracy is consistent with broader evidence that advances in machine learning has substantially improved flood-forecasting (Nearing et al., 2024). Delivery has not kept pace: the rest of this section documents how far the status quo falls short of providing households these forecasts.

5.2 Status Quo Delivery Is Partial

Neither system described in Section 2.2 reliably reaches households. The CWC's warnings only extend to mass media and officials, and the latter is no guarantee that alerts reach households: in a randomized experiment we conducted in this region during the 2019 monsoon, delivering flood alerts to local leaders did not measurably increase local leader or household alert receipt (Figure A4).

Google's notification, reach households directly, but only some of them. Google alerts are pushed to Android phones with location services enabled within the affected area of interest (AOI). To measure reach, we combine Google data on the number of phones notified with settlement-level population counts from Meta's High Resolution Settlement Layer, and treat a household as reached if any of its phones receives the alert.[19] Average reach conditional on issuance was 17.5 percent of AOI households in 2023, 53.1 percent in 2024, and 56.9 percent in 2025 (Figure 6(a)). Household reports tell the same story. In 2024 and 2025, we asked control households whether they had received a flood alert as a phone notification. During the 2024 season, 93 percent lived in a GP with a recorded severe forecast alert and 78 lived in one that flooded, yet only 13 percent reported a phone notification, and 66 percent experienced community flooding without one; 2025 is similar (Figure 6(b)). The gap persists even when stakes are high: water entered

[18] Figure A5 presents all stations for the 2025 monsoon.

[19] Populations are converted to households using the mean control-group household size of 5.48.

the home for 20 percent of control households in 2024 and 15 percent in 2025, almost all in a GP with a recorded severe forecast alert; and three quarters of them received no phone notification. Panel (a) counts alerts Google pushed and panel (b) notifications households recall, so levels are not comparable, but both show the digital channel stopping well short of the households the forecast covered. One reason is awareness: in a phone survey we conducted, roughly 70 percent of smartphone owners were unaware of GPS or location services.

This gap is as stark counting alerts from any source. Figures 7 uses four seasons of household surveys to compare flood exposure, forecast coverage, and self-reported alert receipt among control households. In 2025, 67 percent of households report that their community was flooded and forecasts covered over 90 percent of households, yet fewer than half received a warning from any source; receipt never exceeded half in any season, and in the first two seasons it was roughly one household in five. Among the 15 percent of households whose own home flooded in 2025, virtually all covered by forecasts, fewer than two-thirds were warned. In the pooled survey, 44 percent of control households report experiencing at least one flood with no alert.

Unreliable delivery impacts willingness to act on alerts. At endline, 27 percent of control households report having experienced a false alarm and 32 percent a flood with no warning; among households experiencing either error, 87 percent say it reduced their willingness to act on future alerts. The last-mile problem is therefore both physical and behavioral.

6 Impact on Delivery and Ex-Ante Adaptation

The conceptual framework implies a causal chain: warnings must first reach households', and warned households must act before water arrives. We present results in that order, first on alert delivery Section 6.1 and then on protective action Sections 6.2.

6.1 Treatment Closes the Last-Mile Gap

Table 3 reports ITT effects on alert receipt from equation (1), pooling midline and endline rounds (roughly 20,800 household-round observations, 319 clusters). The first row shows the extensive margin: 35.0 percent of control households received any flood alert and assignment raises this probability by 9.3 percentage, an increase of roughly one quarter. This coefficient is also the first stage for the instrumental-variables estimates that follow. Impacts on the intensive margin are considerably larger: treated households received 4.7 additional alerts over the season against a control mean of 2.1, more than tripling exposure to warnings.

The channel decomposition shows the program added a delivery channel rather than displacing existing ones. Receipt through IT-based channels (SMS, WhatsApp, phone calls) more than

doubles, rising 16.6 percentage points on a control base of 12.8 percent; receipt through community outreach (loudspeakers, flags) rises 1.6 points; and receipt through in-person contact and government mass media (television, radio, newspapers) is unchanged, with precise null coefficients. Because agents relayed the same underlying CWC and Google forecasts, these effects measure delivery, rather than new information production.

These additional warnings arrived earlier and provided more information. The probability receiving an alert before floodwater reached the home rises by 3.2 percentage points on a control base of 11.1 percent, an increase of nearly 30 percent. Treated households are also significantly more likely to report that their alerts contained each element of forecast content the system produces: that a flood was forecast, its expected date and time, water depth in the panchayat, the river's depth and name, the severity level, and the direction of change in the water level, with gains of 28 to 92 percent relative to control means. Neither system's alerts describe the cause of flooding, and treatment effects on cause-related information are correspondingly small: reports of embankment-break information rise 15 percent (significant only at the ten percent level) and dam-release information an insignificant 8 percent. This is consistent with agents relaying alert information rather than adding to it.

Crossing self-reported receipt against self-reported panchayat flooding shows how additional alerts change the structure of forecast errors. Treatment reduces misses, a flood with no alert, by 6.1 percentage points on a control base of 43.9 percent, while modestly increasing false alarms, an alert with no flood, by 3.5 percentage points on a base of 5.2 percent, a mechanical consequence of higher alert volume; the probability of experiencing either error falls by 2.6 percentage points. Perceived quality improves despite the additional false alarms, consistent with misses being costlier: the shares of households that trust alerts completely, find them very useful, and rate them very accurate each rise by roughly 30 percent of the control mean.

Warnings diffuse beyond direct recipients. In the 2025 round, treatment raises the share of households that discussed or shared flood information with anyone else by 10.6 percentage points on a control base of 55 percent, and raises sharing beyond the respondent's family, with residents of the same or other tolas, other villages, people outside the panchayat, or community leaders, by 10.4 percentage points on a base of 54 percent.

6.2 Warnings Shift Action Forward

The midline survey elicited protective steps taken over the monsoon window (Jeshth to Kartik, roughly May to November), which we group into two meta-domains: advance responses, 27 preparatory steps in five domains, taken well ahead of flooding, and last-minute responses taken immediately before or during onset, such as moving household members or evacuating docu-

ments. Each is aggregated into a standardized index following Kling, Liebman and Katz (2007), with the advance-response index an equal-weighted average of the five domain indices.

Table 4 reports that treatment raises the advance-response index by 0.052 standard deviations (significant at 5 percent) and leaves last-minute responses unchanged, with a point estimate near zero (-0.009), the combined index rises by 0.042 standard deviations. Scaling by the first stage, receiving an alert because of the program raises advance preparation by 0.57 standard deviations among compliers, subject to the interpretation caveats in Section **??**. The pattern suggests the program operated as an early-warning technology rather than as a general increase in flood salience: warnings enabled preparation that requires lead time without adding to action at flood onset. The null on last-minute actions also indicates neither panic responses to additional alerts nor substitution away from onset-time protection.

The advance-response effect is broad-based rather than driven by a single margin. Across the five domain sub-indices, treatment raises house and environment preparation by 0.050 standard deviations (significant at 1 percent), food and water storage by 0.041 (10 percent), animal protection by 0.039 (10 percent), and agricultural preparation by 0.031 (10 percent); the human-health sub-index rises by a statistically insignificant 0.028. No domain dominates, consistent with households drawing down a general preparation margin once they know a flood is coming.

The component-level estimates show that responses involved real resource costs. Treated households are 2.3 percentage points more likely to have stocked food (control mean 42.6 percent), 1.5 points more likely to have sandbagged (19.1 percent), and 1.0 points more likely to have dug a drainage canal (8.1 percent). These individually modest margins accumulate across 27 steps into the index effects above (Table A6).

The program delivered more, earlier, richer, and more trusted warnings, and households converted them into anticipatory protection.

7 Impact on Ex-Post Welfare

Section 6 showed that warned households prepared before floods arrived; this section asks whether that preparation reduced welfare losses once water came. Our summary measure is an overall welfare index built from five post-flood domains: physical health, asset damage, agricultural harvest, food security, and coping. Within each domain, variables are standardized against the control group within survey round and aggregated into a Kling, Liebman and Katz (2007) index; the overall index is the equal-weighted average of the six domain indices, oriented so that higher values indicate better welfare. Because the value of a warning should scale with the severity of the flood it warns about, we estimate the dose-response and binned specifications of equations (2) and (3).

7.1 Overall Welfare and the Severity Gradient

We begin with the pooled treatment estimate. Across four seasons, assignment raises the overall index by 0.034 standard deviations (s.e. 0.028), indistinguishable from zero (Table A8, column 1). This is the expected average result for an intervention with state-contingent payoff over a window in which 77 percent of village-years saw no flooding: the pooled coefficient is the incidence-weighted average of a near-zero effect where nothing happened and larger effects where floods arrived. Column 3, which pools the flooded categories, shows the average of the latter: any flooding lowers control-group welfare by 0.226 standard deviations, and treatment offsets 0.133 of that loss, about six-tenths (Table A8, column 3). The table also reports, for completeness, a two-stage least squares estimate that instruments any-alert receipt with assignment (0.373, s.e. 0.329); we do not interpret it, because the exclusion restriction potentially fails for welfare outcomes.

Table 5 unpacks the flooded categories. Flooding is costly in the control group: each additional flood-day lowers welfare by 0.018 standard deviations, and relative to unflooded villages, low and severe flooding reduce welfare by 0.183 and 0.359 standard deviations, respectively. Treatment effects concentrate where these losses occur. In the continuous specification, the treatment-by-severity interaction is 0.015 standard deviations per flood-day, offsetting roughly five-sixths of the marginal welfare cost of an additional day of inundation, we cannot reject complete offset ($p = 0.31$). The binned specification separates the two categories. Under severe flooding, the treatment effect is 0.299 standard deviations against a penalty of 0.359: treatment offsets 83 percent of the loss, and we cannot reject that treated households under severe flooding are as well off as unflooded control households ($p = 0.43$). Under low flooding, the effect is 0.073 against a penalty of 0.183, an offset of 40 percent, and full offset is rejected ($p = 0.009$). The residual losses in the two categories, 0.110 and 0.060 standard deviations, are not statistically distinguishable from each other ($p = 0.47$). We read the pattern as partial protection throughout, more complete under severe flooding in point estimates, with the difference between the categories concentrated in particular domains.

The unflooded village-years provide a falsification test. In village-years with no flooding, the treatment effect is 0.001 standard deviations (s.e. 0.028), and the continuous specification's treatment coefficient at zero flood-days is a similarly null 0.010, consistent with treatment effects operating only through averted flood losses.

Robustness. Because treatment is randomized but severity is not, the treatment-severity interaction could in principle reflect fixed village characteristics correlated with flood exposure, riverbank proximity, elevation, baseline vulnerability, rather than the protective effect of warnings under flooding. Table A5 re-estimates the continuous and binned specifications adding village fixed effects, which absorb all time-invariant village heterogeneity, including the uninteracted

treatment effect; identification then comes only from within-village variation in severity across seasons. The pattern survives this demanding specification: the treatment-severity interaction is 0.010 standard deviations per flood-day (s.e. 0.003) and the severe-flood interaction 0.202 (s.e. 0.070), each significant at the 1 percent level, attenuated by roughly a third relative to the headline estimates of 0.015 and 0.298 $(0.299 - 0.001)$.[20] The offset share is more stable than either coefficient, because measurement error shrinks the severity gradient and its interaction together: within villages, treatment offsets roughly seven-tenths of the per-day welfare gradient (0.010 against 0.014) and 85 percent of the severe-flood penalty (0.202 against 0.238), close to the headline shares, and full offset is again not rejected per flood-day ($p = 0.33$) or under severe flooding ($p = 0.53$), so that protection remains most complete against the worst days. We therefore treat the district-by-round estimates as the headline and the village fixed-effects estimates as confirmation that no time-invariant village attribute can account for the treatment-severity gradient.

The severity measure itself requires two construction choices: the inundation threshold above which a village-day counts as flooded, and the treatment of gaps in the satellite record. Appendix 11.3 shows the results are insensitive to both. Across ten thresholds, from 0.05 to 0.50 (Table A4), the severity gradient is negative and significant at the 1 percent level throughout, and the treatment interaction is positive throughout and significant at the 1 percent level across the interior band from 0.20 to 0.35 containing the paper's definition; replacing the flood-day count with the duration of bridged flood episodes leaves both coefficients significant at the 1 percent level, with the treatment offset largely unaffected.

7.2 Welfare Effects by Domain

Table 6 decomposes the index. Relative to unflooded villages, severe flooding lowers the asset index by 0.370 standard deviations, the harvest index by 0.481, the food-security index by 0.103, and the coping index by 0.174; low flooding lowers the harvest and asset indices by 0.302 and 0.175. Treatment offsets these losses where they arise: under severe flooding, it raises the asset index by 0.217 standard deviations, the harvest index by 0.283, and the food-security index by 0.158; under low flooding, it raises the asset and harvest indices by 0.105 and 0.134 standard deviations. The harvest response runs through the extensive margin of avoiding crop loss: under severe flooding, treatment raises the probability of reporting no crop loss from flooding by 16.3 percentage points against a control-group decline of 19.6 points, while unconditional paddy quantities are too noisy to detect effects of plausible size (Table A7). Physical health is the exception on both margins: the health index shows neither a severity gradient (0.015 and 0.008 in the

[20] Attenuation of this form is expected: the within transformation removes the persistent cross-village component of true severity while retaining measurement error in the satellite-based measure, mechanically amplifying attenuation bias (Griliches and Hausman, 1986).

low and severe bins) nor treatment effects.

The lower panel of Table 6 reports raw outcomes, so beneficial effects carry negative signs. Under severe flooding, treatment reduces the probability that the household's own home flooded by 12.7 percentage points, against a severe-flood increment of 23.7 points: the physical preparation documented earlier, sandbagging, drainage, and securing the house, kept water out of homes. House damage falls by ₹1,907 under severe flooding, offsetting roughly four-fifths of the ₹2,347 severe-flood increment. The probability of food insecurity falls by 6.7 percentage points, more than offsetting the 5.4 point severe-flood increment, consistent with the food stocking induced ex ante.

Flooded control households also finance losses through costly borrowing, imprecisely, by ₹7,762 under severe flooding. Under severe flooding, treatment raises the coping index by 0.211 standard deviations and reduces loan amounts by ₹12,632, roughly 35 percent of the pooled control mean of ₹36,345 and approximately 40 percent of flood-season income. Effects under low flooding are small and insignificant, mirroring the concentration of damages. The is the pattern predicted by the private-preparation channel of the conceptual framework: households that protected assets and stocked essentials ex ante had less damage to finance ex post, and substituted preparation for distress borrowing rather than for other coping margins.

The domain results reproduce the composition the framework predicts. Physical protection is partial in both flood categories: treatment offsets roughly 60 percent of asset losses in each, and 44 and 59 percent of harvest losses, so treated households under severe flooding still carry material losses. What closes completely under severe flooding are the second-order margins: food insecurity and distress borrowing are more than fully offset, while under low flooding, where control households absorb their smaller losses within buffers of savings and stocks, these margins barely move in either arm. Partial protection of assets and harvests was enough to keep treated households inside the buffer that separates a loss absorbed at home from one financed at high interest, which is why the overall index reads as near-complete protection under severe flooding even though physical losses are only partly averted.

7.3 Interaction with State Relief

Early warnings operate alongside a state relief apparatus, raising two questions: does the state respond to floods at all, and does household preparation change what households draw from it? Table 7, using the 2024 and 2025 rounds, answers both. Relief is strongly responsive to severity in the control group: the overall relief index rises by 0.347 and 0.409 standard deviations under low and severe flooding, dry-ration receipt by roughly 7 percentage points, cash transfers of ₹7,000 through Direct Benefit Transfer by 10.4 and 16.4 points, and community-kitchen use by 6.2 and

11.5 points.

Treatment leaves the entitlement components untouched: effects on rations and cash transfers are null in every severity bin, so warned households remained equally likely to receive official transfers, and the program neither crowded out relief nor changed its targeting. The one component that moves is community-kitchen use, which falls by 5.6 percentage points under low flooding and 9.1 points under severe flooding, offsetting most of the flood-driven increase on a pooled control mean of 6.4 percent. Households that stocked food and remained food-secure did not require emergency feeding. Aggregate relief indices show correspondingly negative but imprecise treatment interactions. These averted relief costs accrue to the state and are excluded from the benefit side of the cost-effectiveness calculation, one of the margins on which it is conservative.

In sum, the welfare effects of last-mile delivery are concentrated in flooded village-years, increasing in severity, largest where floods are worst, and traceable item by item to the preparation the warnings induced.

8 Citizen Engagement, Valuation and Cost Effectiveness

This section examines households' perceptions of the system's reliability and their worry about future flooding, their deeper beliefs about science and technology, and their revealed willingness to pay for alert delivery and, finally, whether last-mile delivery is worth its cost.

8.1 Reduced Concern About Unexpected Flooding

At the 2025 endline, households reported whether, over the past few years, they had received false alerts or failed to receive alerts, and how worried they were about unexpected flooding in 2026. Table 8 relates these perceptions to treatment and to cumulative flood exposure over the four study seasons, measured as the village's average flood-days per year and as the number of seasons in which the village flooded at all. Exposure counts any flooding, low or severe, since beliefs form over the full history of events against which warnings could be checked.

Perceived reliability improves on the error margin that matters. Treatment reduces the share of households reporting missed alerts by 4.8 percentage points on a control base of 31.7 percent, and by a further 1.8 points per average flood-day. Looking across frequency bins, reported misses rise with exposure in the control group: by 13.5 and 11.6 percentage points among households in villages flooded in two and in three or four of the four seasons (5.1 points, insignificant, where the village flooded once). Treatment closes this gap. Perceived false alarms are unchanged on average (0.3 points, insignificant, on a base of 27.5 percent) and decline with exposure under

treatment: 0.8 points per average flood-day, and 5.5 points, marginally, in the most frequently flooded category. Although measured false alarms rose mechanically with alert volume (Section 6.1), treated households do not perceive the system negatively, consistent with them weighting avoided misses more heavily.

Worry shows the same pattern. Flood exposure raises concern about future flooding in the control group, by 1.0 percentage point per average flood-day, from a base of 78.3 percent; treatment offsets roughly nine-tenths of the per-day increase. Across all three outcomes, treatment effects among never-flooded households are near zero: treatment lowers concern of being caught unwarned, not concern about floods in general. This benefit is unpriced in the cost-effectiveness calculation.

8.2 Belief in Science and Technology

We next ask whether the program changed households' beliefs about science and technology. The 2024 and 2025 endlines asked whether the season's flooding experience had changed the household's perceptions of science and technology, and fielded a four-item sentiment battery adapted from the World Values Survey module. Table 9 reports the binned specification of equation (3).

Among treated households in severely flooded villages, the probability of reporting that flooding changed their views of science and technology a lot rises by 10.6 percentage points on a control base of 43.6 percent, and the sentiment index rises by 0.117 standard deviations, with components moving toward science. Treated households under severe flooding are 6.5 percentage points more likely to strongly reject the statement that science is unimportant in daily life and 3.3 points more likely to strongly reject the statement that science erodes right and wrong. Every interaction in the no-flood bin is a precise zero. Households revise their views only when they both receive the warnings and experience a severe flood. Severe flooding alone also raises agreement that the world is better off because of science and technology by 6.2 percentage points among control households, possibly suggesting some updating from exposure to forecasts through status-quo channels.

8.3 Willingness to Pay for Alert Delivery

The 2025 endline elicited willingness to pay using an incentivized Becker–DeGroot–Marschak mechanism with descending price lists. Households completed three elicitations: a non-incentivized practice round for a notebook, an incentivized round for a monsoon-onset forecast message, and an incentivized round for a flood-onset alert service, a message sent by a member of Yuganter, the implementing NGO, whenever a flood is forecasted in the household's area, with a maximum price of ₹100 in the incentivized rounds. Treated households were told this was the same service

they had been receiving. Following Dizon-Ross and Jayachandran (2022), the practice and monsoon elicitations serve as household-specific benchmarks that absorb elicitation noise and generic propensity to bid.

Table 10 shows a valuation premium specific to the experienced service. Treatment has no effect on practice-round bids or on bids for the monsoon-onset forecast. Willingness to pay for flood alert delivery, by contrast, is ₹1.74 higher among treated households conditional on both benchmarks (significant at 5 percent), about 5 percent of the unconditional control mean of ₹34; rescaled by the first stage, households induced to receive alerts value the service ₹12.5 higher, roughly 37 percent of the control mean. Because bids were capped at ₹100, measured means understate underlying valuation and the treatment premium is, if anything, attenuated, so we read these estimates as evidence on direction and incidence rather than as the full magnitude of demand.

Positive valuation does not require flood experience: the service is insurance, so any household that assigns positive probability to future flooding and some skill to the forecast should bid a positive amount. Experience changes is knowledge of the service's quality, which is why the premium loads on treated households while the level of demand need not depend on recent exposure. Demand is far from universal, with 52.5 percent of control households bidding a positive amount and a conditional mean of ₹65, but it is economically meaningful: control-group willingness to pay of ₹34 per season is roughly one and a half times the program's delivery cost of ₹23 per household-season, a comparison the ₹100 cap can only have made harder to pass.

The monsoon null also helps interpret Section 8.2: treatment raised sentiment toward science and demand for the service households has experienced, while an unverified product in the same category commanded no premium, pointing to experience-based learning. Together, the demand-side results address the trust issue documented earlier: under status-quo dissemination, forecast errors erode willingness to act on warnings; under last-mile delivery, households report fewer misses, worry less, update toward the technology, and pay more for the service.

8.4 Cost-Effectiveness

Households' willingness to pay exceeds the program's delivery cost; we now make the comparison of benefits and costs directly. The calculation, summarized in Table 11, uses one experimentally estimated benefit channel on the numerator and the program's measured delivery costs on the denominator. It is deliberately conservative, and we report the margins on which it is so.

The benefit side uses a single channel: averted distress borrowing under severe flooding. Households finance uninsured flood losses largely through credit, so reduced seasonal borrowing provides a reasonably comprehensive money-metric summary of averted out-of-pocket losses.

Restricting the numerator to this channel also avoids double counting margins that borrowing finances, such as house repair. From Table 6, treatment reduces borrowing by ₹12,632 per household under severe flooding; effects under low flooding are small and excluded. Over 2022 to 2025, the probability that a covered household experienced severe flooding was 5.4 percent per season, so the expected benefit is ₹682 per covered household-season. The program covered roughly 320,000 households (160 treatment panchayats with about 2,000 households each) over four seasons, for total benefits of ₹873.1 million.

The cost side counts measured delivery costs only: agent-based dissemination cost ₹46,509 per panchayat-season, or ₹29.8 million over 640 panchayat-seasons, ₹23 per household-season. We take forecast production by the CWC and Google as given: both systems operate at scale regardless of the program, so the policy-relevant margin is the incremental cost of delivering existing forecasts to households.

The resulting benefit–cost ratio is 29:1, and it is conservative on two margins. First, the numerator excludes every other measured gain: asset protection (0.22 standard deviations), harvests (0.15), food security (0.16), the 12.7 percentage point reduction in home flooding, averted public relief spending including the decline in community-kitchen use, and reduced worry. Second, the study window was mild: severe-flood incidence of 5.4 percent over 2022 to 2025 compares with 15.7 percent over 2017 to 2025, and at the historical rate the ratio exceeds 85:1.

A comparison with the state's relief spending put the cost in additional context: ₹23 per household-season is roughly 2 percent of the expected cost of the cash transfer alone, approximately INR 1,100 per household-season at historical severe-flood incidence (₹7,000 per Direct Benefit Transfer times 15.7 percent), before counting rations, supplies, evacuation, and rehabilitation. Delivering accurate forecasts the last mile is inexpensive relative to both what households will pay for it and what the state spends responding to floods.

9 Conclusion

Flood forecasting has made remarkable strides. Forecasts for the rivers of north Bihar are accurate at lead times of several days, produced at scale by a government agency and a technology company, and freely available. This paper shows that their value is realized only when warnings are received, trusted, and actionable. We developed and tested a delivery layer built on community engagement and training: local agents relayed official warnings, unmodified, through loudspeakers, flags, SMS, and WhatsApp, at a cost of ₹23 per household-season. The program raised alert receipt by nine percentage points, more than tripled the number of warnings households received each season, and induced earlier and better preparation.

The treatment paid off in the states of the world it was designed for. Under severe flooding,

treated households were less likely to have floodwater enter their homes, suffered less house damage, were less likely to be food insecure, borrowed roughly INR 12,600 less, and recovered most of the welfare penalty that severe floods impose on control households, an offset statistically indistinguishable from complete. Village-years without floods showed no effect. Households that experienced the service became less likely to worry about unexpected future flooding and updated toward the science where floods verified it, and were willing to pay for delivery at levels above the program's cost.

Three lessons follow. First, prediction and dissemination are different goods, and protection requires both. Second, the value of a warning system is state-contingent, and evaluations pooled over short windows will understate it. Payoffs arrive only when floods do, and short study periods samples severe floods sparsely. Reweighting our estimates to the region's 2017 to 2025 severe-flood incidence raises the benefit-cost ratio from 29 to 1 to more than 85 to 1. Third, households financed their own preparation: the gains required no accompanying transfers, and treated households drew no more government relief than control households. They also recognized the value of delivery. Willingness to pay exceeds the ₹23 delivery cost even among control households, and experience with the service raises it further. What warnings could achieve jointly with anticipatory action, such as forecast-triggered cash transfers, is a natural next experiment.

Two study design limitation suggest an agenda for future work. The experiment prices a delivery bundle: local agents supplied both reach and a trusted messenger. Whether an automated system speaking in a trusted local voice, or the same reach behind an unfamiliar sender, would achieve these effects determines how cheaply and how far the model could scale. And the point extends beyond floods: artificial intelligence is making prediction feasible across hazards where households in low-income settings bear uninsured risk, from crop disease to heat to epidemic onset. Each hazard differs in timescale, footprint, and the protective actions it admits, so the dissemination layer that turns its forecasts into action will differ too. In north Bihar, delivering a season of flood alerts to a household cost less than a cup of tea a month, and it turned prediction into protection.

References

Agrawal, Ajay, Joshua Gans, and Avi Goldfarb. 2018. *Prediction Machines: The Simple Economics of Artificial Intelligence.* Boston, MA:Harvard Business Review Press.

Balboni, Clare, Oriana Bandiera, Robin Burgess, Anton Heil, Clément Mazet-Sonilhac, Munshi Sulaiman, and Yifan Wang. 2026. "Weathering Poverty." London School of Economics Working Paper.

Bandyopadhyay, Somnath. 2017. "Natural Floods and Unnatural Disasters in Bihar." *Dialogue Earth* (*formerly China Dialogue*), Published 30 March 2017.

Björkman Nyqvist, Martina, Andrea Guariso, Jakob Svensson, and David Yanagizawa-Drott. 2019. "Reducing Child Mortality in the Last Mile: Experimental Evidence on Community Health Promoters in Uganda." *American Economic Journal: Applied Economics*, 11(3): 155–192.

Burlig, Fiona, Amir Jina, Erin M. Kelley, Gregory Lane, and Harshil Sahai. 2024. "Long-Range Forecasts as Climate Adaptation: Experimental Evidence from Developing-Country Agriculture." National Bureau of Economic Research NBER Working Paper 29721.

Camacho, Adriana, and Emily Conover. 2019. "The Impact of Receiving SMS Price and Weather Information in Colombia's Agricultural Sector." *Food Policy*, 83: 161–174.

Carleton, Tamma, Esther Duflo, B. Kelsey Jack, and Guglielmo Zappalà. 2024. "Adaptation to Climate Change." In *Handbook of the Economics of Climate Change*. Vol. 1, , ed. Lint Barrage and Solomon Hsiang, Chapter 4, 143–248. Amsterdam:North-Holland.

Central Water Commission. 2020. "Flood Forecasting / Hydrological Observation." *Central Water Commission, Ministry of Jal Shakti, Department of Water Resources, River Development and Ganga Rejuvenation, Government of India*.

Central Water Commission. 2024. "Flood Forecasting and FloodWatch India." *Central Water Commission, Ministry of Jal Shakti, Government of India*, Over 330 river-monitoring stations; approximately 10,000 forecasts issued annually. Accessed June 2026.

Christian, Paul J., Felipe A. Dunsch, Jonas Heirman, Erin M. Kelley, Florence Kondylis, Gregory Lane, Jennifer Waidler, Nidhila Adusumalli, Odbayar Batmunkh, and Kriti Malhotra. 2025. "Fast Action for Floods: RCT Evidence on Forecast-based Cash Transfers from Bangladesh and Nepal." National Bureau of Economic Research NBER Working Paper 34414.

Coughlan de Perez, Erin, Kristoffer B. Berse, Lianne Angelico C. Depante, Evan Easton-Calabria, Elton Pierre R. Evidente, Theodore Ezike, et al. 2022. "Learning from the Past in Moving to the Future: Invest in Communication and Response to Weather Early Warnings to Reduce Death and Damage." *Climate Risk Management*, 38: 100461.

Cristia, Julian, Pablo Ibarrarán, Santiago Cueto, Ana Santiago, and Eugenio Severín. 2017. "Technology and Child Development: Evidence from the One Laptop per Child Program." *American Economic Journal: Applied Economics*, 9(3): 295–320.

District Nalanda, Government of Bihar. n.d.. "Disaster Management: Bihar Hazard Profile." *District Administration, Nalanda, Government of Bihar*, States that 22.1% of India's total flood-affected population is located in Bihar, and that annual flooding in Bihar accounts for about 30–40% of India's flood damages.

Dizon-Ross, Rebecca, and Seema Jayachandran. 2022. "Improving Willingness-to-Pay Elicitation by Including a Benchmark Good." *AEA Papers and Proceedings*, 112: 551–555.

Dodge, Eric, Yusuf Neggers, Rohini Pande, and Charity Troyer Moore. 2025. "From Delay to Payday: Easing Bureaucrat Access to Implementation Information Strengthens Social Protection Delivery." National Bureau of Economic Research Working Paper 33756.

Dupas, Pascaline. 2014. "Short-Run Subsidies and Long-Run Adoption of New Health Products: Evidence from a Field Experiment." *Econometrica*, 82(1): 197–228.

Fabregas, Raissa, Michael Kremer, and Frank Schilbach. 2019. "Realizing the Potential of Digital Development: The Case of Agricultural Advice." *Science*, 366(6471): eaay3038.

Fafchamps, Marcel, and Bart Minten. 2012. "Impact of SMS-Based Agricultural Information on Indian Farmers." *The World Bank Economic Review*, 26(3): 383–414.

Global Commission on Adaptation. 2019. "Adapt Now: A Global Call for Leadership on Climate Resilience." Global Commission on Adaptation and World Resources Institute, Rotterdam and Washington, DC. Estimates that early warning systems save lives and assets worth at least ten times their cost; US$800 million spent on such systems in developing countries would avoid losses of US$3–16 billion per year. Accessed June 2026.

Google Research. 2019. "An Inside Look at Flood Forecasting." *Google Research Blog*, Published 19 September 2019.

Google Research. 2026. "Flood Forecasting." *https://sites.research.google/gr/floodforecasting/*, Accessed September 8, 2026.

Griliches, Zvi, and Jerry A. Hausman. 1986. "Errors in Variables in Panel Data." *Journal of Econometrics*, 31(1): 93–118.

Gupta, Vineet, Maulik Jagnani, Miray Omurtak, Rohini Pande, and George Taylor. 2026. "Empirical Flood-Damage Functions for Agriculture: Duration, Experience, and Recovery." Working paper.

Hanna, Rema, Esther Duflo, and Michael Greenstone. 2016. "Up in Smoke: The Influence of Household Behavior on the Long-Run Impact of Improved Cooking Stoves." *American Economic Journal: Economic Policy*, 8(1): 80–114.

Kling, Jeffrey R., Jeffrey B. Liebman, and Lawrence F. Katz. 2007. "Experimental Analysis of Neighborhood Effects." *Econometrica*, 75(1): 83–119.

Lane, Gregory. 2024. "Adapting to Climate Risk With Guaranteed Credit: Evidence From Bangladesh." *Econometrica*, 92(2): 355–386.

Lang, Megan, Jonah Rexer, Siddharth Sharma, and Margaret Triyana, ed. 2025. *From Risk to Resilience: Helping People and Firms Adapt in South Asia. South Asia Development Matters*, Washington, DC:World Bank.

Li, Zhi, Lorenzo Rosa, and Steven Gorelick. 2025. "Severe Floods Significantly Reduce Global Rice Yields." *Science Advances*, 11: eadx7799.

Lloyd's Register Foundation. 2024. "World Risk Poll 2024: Resilience in a Changing World." Lloyd's Register Foundation, London, UK. Gallup survey of approximately 147,000 people across 142 countries; 2023 data wave. 30% of disaster-affected respondents received no warning.

Merz, B., H. Kreibich, R. Schwarze, and A. Thieken. 2010. "Review Article "Assessment of Economic Flood Damage"." *Natural Hazards and Earth System Sciences*, 10(8): 1697–1724.

Molina, Renato, and Ivan Rudik. 2024. "The Social Value of Hurricane Forecasts." National Bureau of Economic Research NBER Working Paper.

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

National Disaster Management Authority, and Centre for Development of Telematics. 2024. "SACHET: Integrated Alert System." *National Disaster Management Authority and C-DOT, Government of India*, CAP-based national public-alert platform; location-based SMS in 19+ languages. Accessed June 2026.

National Remote Sensing Centre. 2020. "Flood Hazard Atlas of Bihar." National Remote Sensing Centre, Indian Space Research Organisation, Department of Space, Government of India, Hyderabad, India. Version date March 2020; published July 2020.

Nearing, Grey, Deborah Cohen, Vusumuzi Dube, Martin Gauch, Oren Gilon, Shaun Harrigan, Avinatan Hassidim, Daniel Klotz, Frederik Kratzert, Asher Metzger, Sella Nevo, Florian Pappenberger, Christel Prudhomme, Guy Shalev, Shlomo Shenzis, Tadele Yednkachw Tekalign, Dana Weitzner, and Yossi Matias. 2024. "Global Prediction of Extreme Floods in Ungauged Watersheds." *Nature*, 627(8004): 559–563.

NITI Aayog. 2023. "National Multidimensional Poverty Index: A Progress Review 2023." NITI Aayog, Government of India, New Delhi. Ranks Bihar as India's poorest state, with the highest multidimensional poverty headcount ratio.

Page, Lucy, and Rohini Pande. 2018. "Ending Global Poverty: Why Money Isn't Enough." *Journal of Economic Perspectives*, 32(4): 173–200.

Pappenberger, F., H. L. Cloke, D. J. Parker, F. Wetterhall, D. S. Richardson, and J. Thielen. 2015. "The Monetary Benefit of Early Flood Warnings in Europe." *Environmental Science & Policy*, 51: 278–291.

Patel, Dev. 2024. "Floods: Measurement, Impacts, and Adaptation." Working paper.

Patel, Dev. 2025. "Environmental Beliefs and Adaptation to Climate Change." Working paper.

Pople, Ashley, Ruth Hill, Stefan Dercon, and Ben Brunckhorst. 2026. "The Importance of Being Early? Anticipatory Cash Transfers for Flood-Affected Households." *The World Bank Economic Review*. Advance article, e-location ID lhag016.

Prasad, Eklavya, Saanjali Verma, and Siddharth Patil. 2025. "Household-level Flood Loss Assessment 2024: 2,290 Households, 134 Wards, 21 Panchayats, Seven Districts, North Bihar." Megh Pyne Abhiyan Field Assessment Report, NOIDA, India. Supported by Tata Trusts. Purposive sample of most-affected wards following the 2024 floods; figures are descriptive and not population representative.

Press Information Bureau. 2022. "Flood Forecasting and Early Warning System." *Press Information Bureau, Ministry of Jal Shakti, Government of India*, PRID 1849954. Accessed June 2026.

Press Information Bureau. 2023. "Flood Forecasting and Management." *Press Information Bureau, Ministry of Jal Shakti, Government of India*, PRID 1946495. Accessed June 2026.

Reichstein, Markus, Vitus Benson, Jan Blunk, Gustau Camps-Valls, Felix Creutzig, Carina J. Fearnley, Boran Han, Kai Kornhuber, Nasim Rahaman, Bernhard Schölkopf, et al. 2025. "Early Warning of Complex Climate Risk with Integrated Artificial Intelligence." *Nature Communications*, 16: 2564.

Rentschler, Jun, Melda Salhab, and Bramka Arga Jafino. 2022. "Flood Exposure and Poverty in 188 Countries." *Nature Communications*, 13(1): 3527.

Rosenzweig, Mark R., and Christopher R. Udry. 2013. "Forecasting Profitability." National Bureau of Economic Research NBER Working Paper 19334.

Rosenzweig, Mark R., and Christopher Udry. 2020. "External Validity in a Stochastic World: Evidence from Low-Income Countries." *The Review of Economic Studies*, 87(1): 343–381.

Shrader, Jeffrey G., Laura Bakkensen, and Derek Lemoine. 2023. "Fatal Errors: The Mortality Value of Accurate Weather Forecasts." National Bureau of Economic Research NBER Working Paper 31361.

Teisberg, Thomas J., and Rodney F. Weiher. 2009. "Background Paper on the Benefits and Costs of Early Warning Systems for Major Natural Hazards." The World Bank Global Assessment Report on Disaster Risk Reduction, background paper, Washington, DC. Estimates EWS benefits for South Asia exceeding costs by more than 500:1 over a decade of normal floods.

United Nations Office for Disaster Risk Reduction. 2025. "Global Assessment Report on Disaster Risk Reduction 2025 — Hazard Exploration: Floods." United Nations Office for Disaster Risk Reduction (UNDRR), Geneva, Switzerland.

United Nations Office for Disaster Risk Reduction, and World Meteorological Organization. 2024. "Global Status of Multi-Hazard Early Warning Systems 2024." UNDRR and WMO, Geneva, Switzerland.

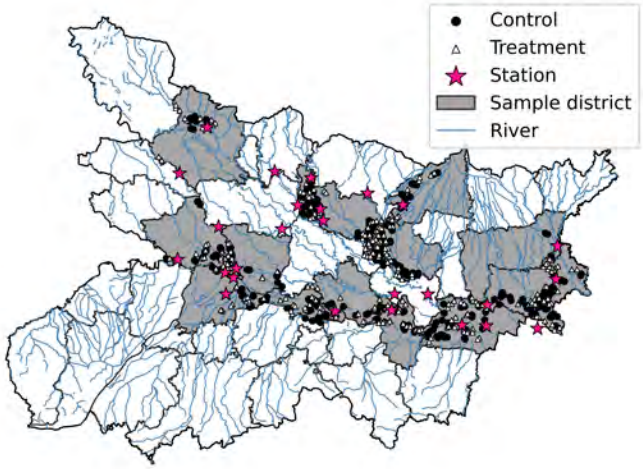
Water Resources Department, Government of Bihar. n.d.. "History of Flood in Bihar." *Flood Management Information System Cell (FMISC), Water Resources Department, Government of Bihar*, States Bihar is India's most flood-prone state, with 76% of North Bihar's population under recurring flood threat and 73.06% of its area flood-affected.

World Meteorological Organization. 2022. "Early Warnings for All: Executive Action Plan 2023–2027." *World Meteorological Organization*, UN/WMO initiative; proposes a US$3.1 billion investment to protect everyone on Earth with early warning systems by 2027. Accessed June 2026.

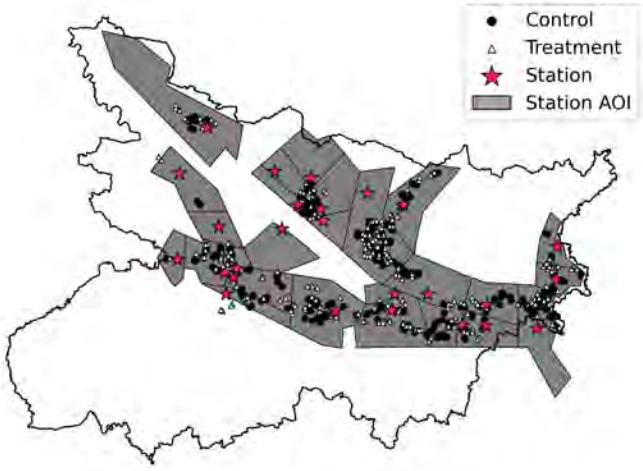
Yegbemey, Rosaine N., Gunther Bensch, and Colin Vance. 2023. "Weather Information and Agricultural Outcomes: Evidence from a Pilot Field Experiment in Benin." *World Development*, 167: 106178.

10 Figures and Tables

Figure 1. Study area: sample locations, rivers, and CWC station areas of interest



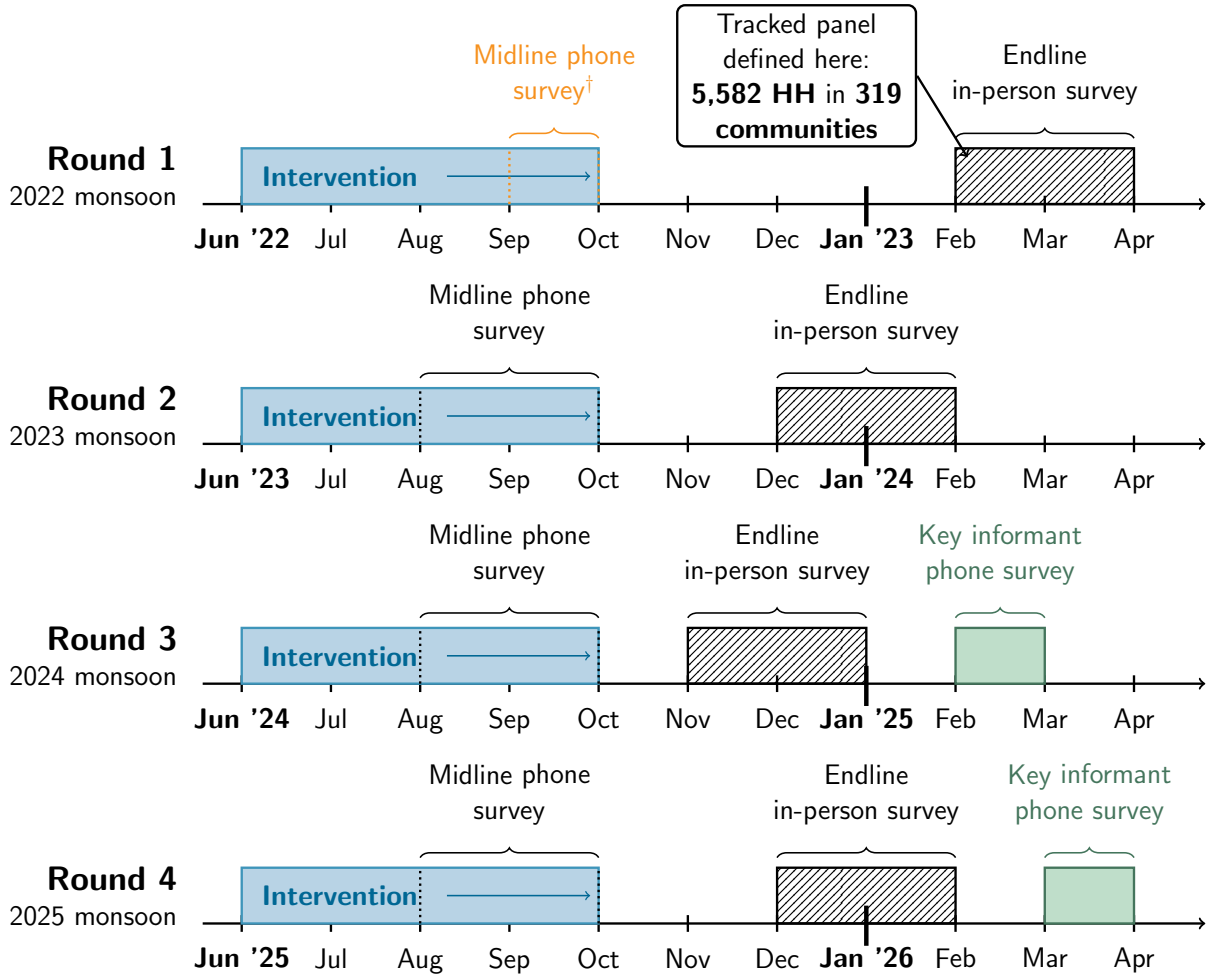
(a) North Bihar: rivers, gauge stations, 319 GPs by arm



(b) CWC station areas of interest and sample locations

Notes: Panel (a) maps the study area in Bihar. Black circles and open triangles mark the 319 sample gram panchayats assigned to the control and treatment arms, respectively; pink stars mark Central Water Commission (CWC) gauge stations; blue lines show rivers; and gray shading identifies sample districts. Panel (b) overlays the same sample locations and gauge stations on the CWC stations' areas of interest (AOIs), shown in gray.

Figure 2. Timeline: randomization, survey rounds, four monsoons



Intervention. Flood-season alert delivery, June–October.

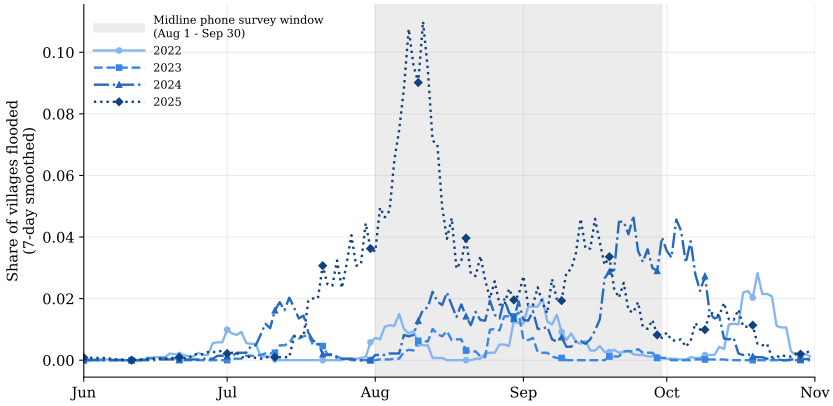
Midline phone survey. Ex-ante outcomes: alert receipt, trust, timing of protective actions.

Endline in-person survey. Ex-post outcomes: exposure, assets, health, harvest, food security, coping.

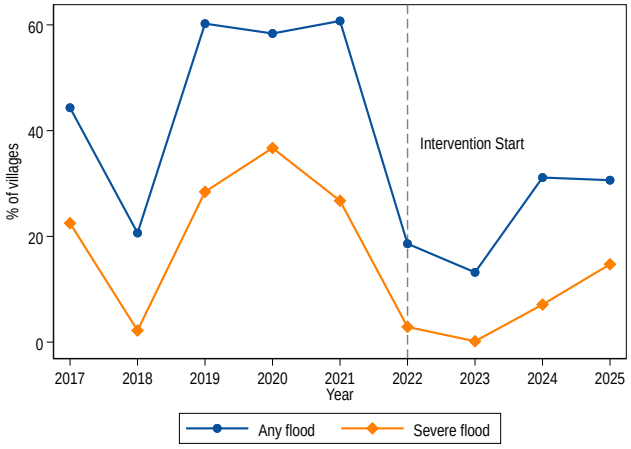
Key informant phone survey. Community-level actions; mukhiyas and ward members across 319 GPs.

Notes: The figure summarizes intervention and survey timing over four monsoon seasons. Randomization was conducted at the gram-panchayat level. The tracked panel was defined at the Round-1 (2022) endline and comprised 5,582 households in 319 communities; these households were re-surveyed at each annual endline. Flood-season alerts were delivered from June through October in each intervention year. Midline phone surveys measured first stage outcomes, including alert receipt, trust, as well as protective actions undertaken by households; from 2023 onward they used the tracked household panel, while the 2022 midline survey (†) used a key-informant-referred resident sample. Endline in-person surveys measured ex-post impacts, assets, health, harvest, food security, and coping. Key-informant phone surveys of mukhiyas and ward members measured community-level outcomes after the 2024 and 2025 endlines only.

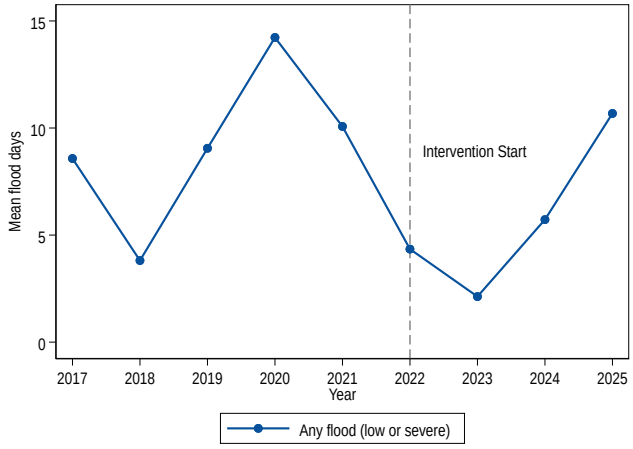
Figure 3. Flood measure: within-season timing, incidence, and duration



(a) Within-season timing of flooding, four monsoons



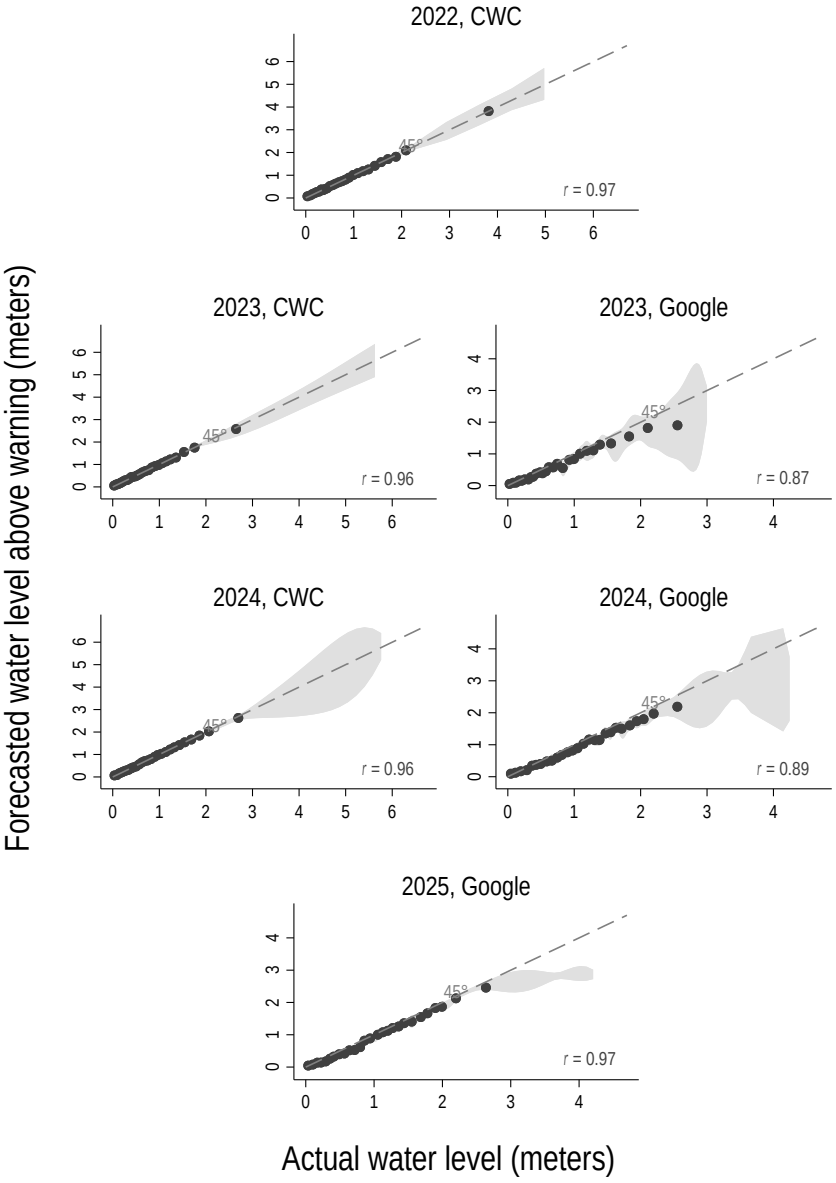
(b) Satellite flood measure, 2017–2025



(c) Flood days conditional on flooding, 2017–2025

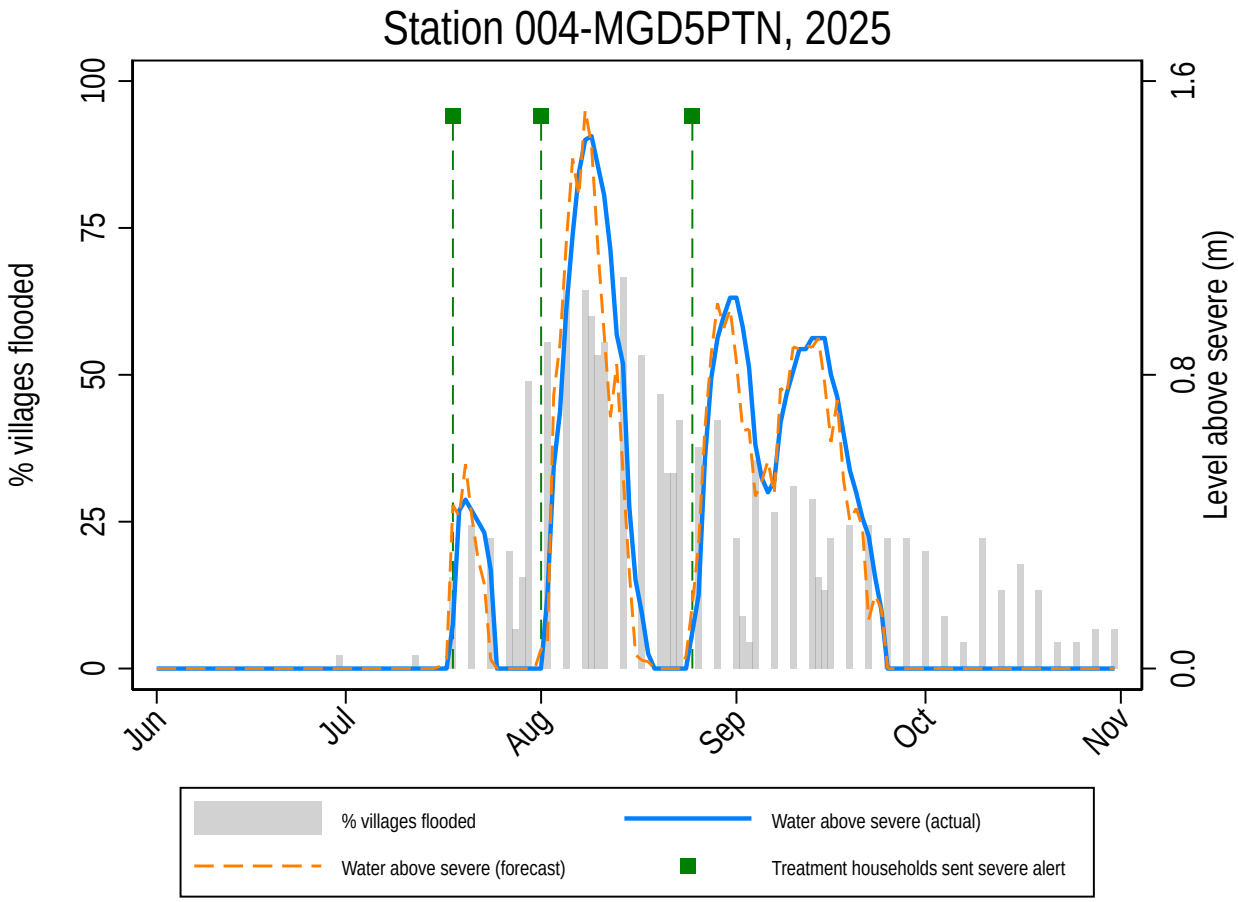
Notes: Panel (a) plots the daily share of sample villages classified as flooded by the synthetic-aperture-radar (SAR) measure during the 2022–2025 monsoon seasons. Series are aligned by days since the first observation in each year and shown as centered seven-day moving averages; the shaded region marks the August 1–September 30 midline phone-survey window. Panel (b) reports, for each year from 2017 through 2025, the percentage of villages with any flooding (the low- or severe-flood categories) and with severe flooding. Panel (c) reports mean SAR flood days among villages classified as having any flood in that year. The dashed vertical line marks the 2022 start of the intervention.

Figure 4. Forecasted vs. realized water levels, by year and source



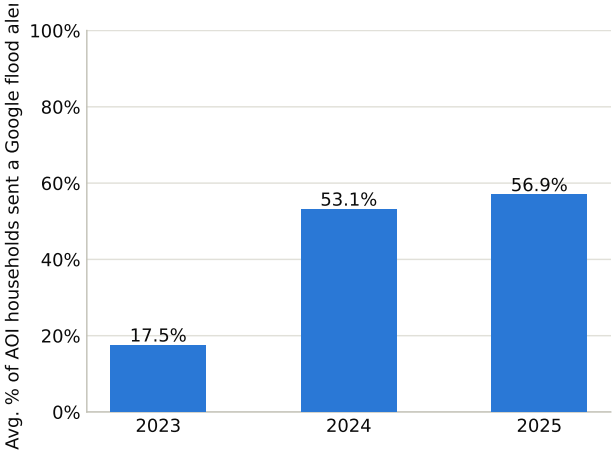
Notes: Each panel compares forecasted and realized station-day water levels for observations in which the realized level exceeded the station-specific warning level. Both axes measure meters above that warning level, with negative exceedances recoded to zero. Forecasts are matched to the date for which the forecast was issued, rather than the bulletin date; where more than one forecast is available for a station-date, the highest forecast is retained. For CWC data, the highest realized exceedance is also retained when multiple bulletins report the same station-date. Dots are binned estimates based on 27 quantile-spaced bins, gray shading is the confidence band from the binsreg fit, and the dashed 45-degree line denotes an equal relationship between forecasted and realized levels. The reported r is the Pearson correlation in the plotted sample. Google data are unavailable for 2022, and CWC data for 2025 are omitted because bulletin coverage ends midyear. Sources are CWC daily water-level bulletins and Google FloodHub data.

Figure 5. Validation: satellite measure vs. water levels vs. alerts, one station

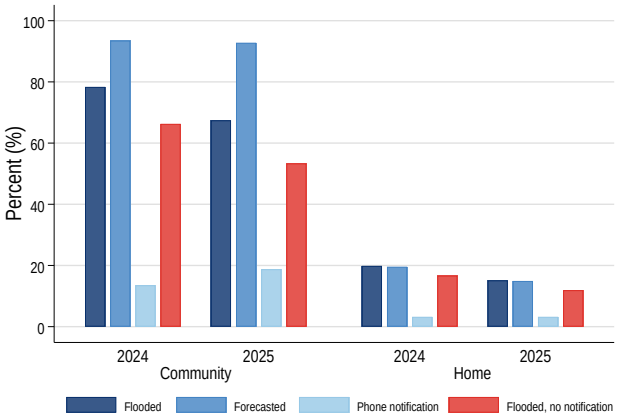


Notes: The figure compares daily flood exposure, water levels, and alert timing for station 004-MGD5PTN during 2025. Gray bars, measured on the left axis, report the percentage of sample villages mapped to the station that are classified as flooded using synthetic-aperture-radar (SAR) data. The solid blue and dashed orange series, measured on the right axis, report realized and forecasted water levels, respectively, in meters above the station-specific severe threshold. Dashed green vertical lines capped by squares mark the start dates of severe-alert events for treatment households.

Figure 6. The digital channel: notifications sent and notifications received



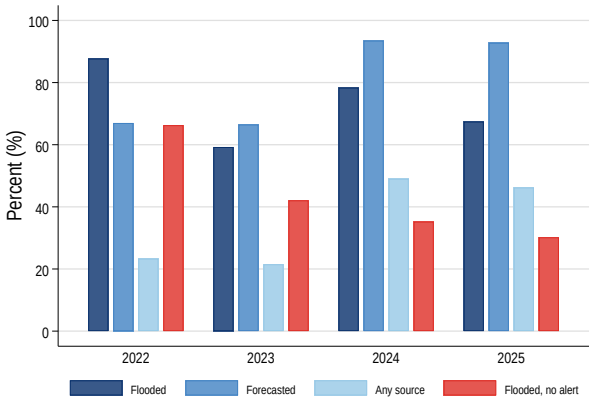
(a) Sent, Google administrative data



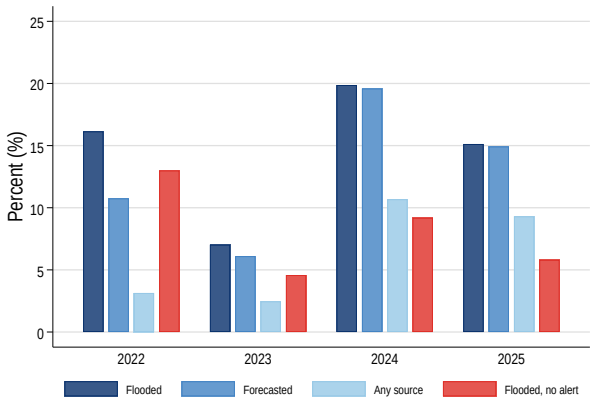
(b) Received, household reports

Notes: Panel (a) reports the unweighted annual mean across station areas of interest (AOIs) of the estimated share of households sent a Google flood alert, conditional on issuance. Sends are the maximum recorded within each station-year; AOI population is measured using the High Resolution Settlement Layer and converted to households using 5.479 persons per household, and the calculation is restricted to station-years represented in the alert files. Panel (b) uses control-household survey responses from the 2024 and 2025 rounds, the only rounds in which the phone-notification item was fielded; the item is asked of control households alone and reads, "During the [year] monsoon season, did you receive any flood alert on your phone through notifications?", with surveyors instructed that it excludes SMS and WhatsApp messages. Within each year, the Community group reports the share of control households who report that their gram panchayat was flooded flooded, the share living in a gram panchayat with a recorded severe forecast alert, the share reporting receipt of a flood alert as a phone notification, and the share reporting community flooding and no phone notification. The Home group reports the corresponding joint shares for flooding that entered the home—not rates conditional on home flooding.

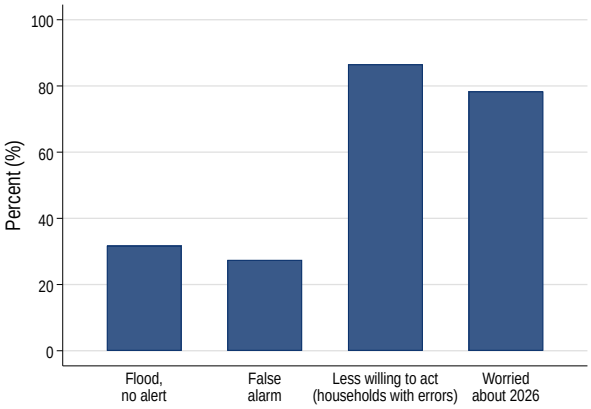
Figure 7. Alert receipt and perceived reliability



(a) Community flooded



(b) Home flooded



(c) Perceived alert errors and worry

Notes: Panels (a) and (b) report unweighted household-level means by year using control-household observations from 2022–2025. In panel (a), Flooded is the share answering yes when asked whether flood waters entered the respondent's gram panchayat during that year's survey reference period; Forecasted is the share living in a gram panchayat with a recorded severe forecast alert; Any source is the share answering yes when asked whether the respondent received any flood warning or alert; and Flooded, no alert is the share reporting community flooding and no receipt of any warning or alert. In panel (b), Flooded is the share reporting that flood waters entered the home during the survey reference period. Forecasted, Any source, and Flooded, no alert are joint shares that report home flooding and, respectively, live in a gram panchayat with a recorded severe forecast alert, report receipt of any flood warning or alert, or report no such receipt. Panel (c) reports unweighted means for control households in the 2025 endline. Flood, no alert is the share answering yes to "In the past few years, did you experience floods but did not receive an alert about it?" and False alarm is the share answering yes to "In the past few years, did you receive any alerts that turned out to be false alarms (i.e., flood water did not arrive as predicted)?". Less willing to act comes from "How much did that affect your willingness to act on future alerts?", which is asked only of households reporting a missed flood or a false alarm, so its denominator is the subset of control households experiencing either error rather than all control households. Worried about 2026 is the share with any concern about "How worried are you about unexpected flooding in 2026?", asked on a five-point scale running from not at all concerned to extremely concerned.

Table 1. Summary statistics: control households, 2022

	Mean	SD	N
Household characteristics (2022 baseline)			
# Males	2.805	1.408	2744
# Females	2.676	1.528	2744
# Under 18s	2.378	1.837	2744
HoH age	44.694	14.913	2744
Mean age	30.274	12.106	2744
HH member 65+ (=1)	0.404	0.643	2744
HoH years of education	8.798	3.499	2744
Muslim (=1)	0.099	0.299	2727
SC/ST (=1)	0.214	0.410	2719
OBC/EBC (=1)	0.593	0.491	2719
Smartphone (=1)	0.915	0.278	2744
Roof/terrace (=1)	0.544	0.498	2724
Land owned (acres)	0.985	1.631	2742
Flood-season income (INR, thousands)	31.084	41.570	2700
Distance to nearest river (km)	2.755	2.162	2744
Agricultural occupation (=1)	0.677	0.468	2732
Agricultural occupation (on farm) (=1)	0.550	0.498	2732
# Rooms in house	3.280	1.936	2619
Pakka House (=1)	0.410	0.492	2619
Use toilet (=1)	0.723	0.448	2739
# Years flooded (2017–2021)	1.767	1.292	2744
Households			2744
Panchayats			159

Notes: The table reports unweighted means, sample standard deviations, and non-missing observation counts for control households observed in the 2022 baseline round. Household composition, household-head characteristics, religion and caste, smartphone access, housing, land, occupation, toilet use, and flood-season income are measured from the 2022 household record; flood-season income is reported in thousands of Indian rupees. The pre-period flood-history measure counts the number of years flooded from 2017 through 2021, and distance to the nearest river is measured in kilometers. The final two rows report the numbers of households and gram panchayats in the sample. Indicators labeled (=1) are binary. Observation counts are variable specific because of item nonresponse.

Table 2. Baseline balance, treatment vs. control

	Control Mean [SD]	Treatment Mean [SD]	Difference	p-value	N
Time-invariant characteristics (imputed from 2022 baseline)					
# Males	2.815 [1.410]	2.803 [1.489]	-0.013	0.785	20890
# Females	2.681 [1.528]	2.650 [1.587]	-0.029	0.521	20890
# Under 18s	2.395 [1.842]	2.393 [1.920]	0.001	0.985	20890
HoH age	44.915 [14.927]	45.044 [15.342]	0.095	0.844	20890
Mean age	30.296 [12.062]	30.843 [12.547]	0.519	0.232	20890
# Members 65+	0.408 [0.647]	0.448 [0.674]	0.039	0.038	20890
HoH years of education	8.778 [3.502]	8.811 [3.443]	0.021	0.876	20886
Muslim (=1)	0.096 [0.295]	0.074 [0.261]	-0.020	0.268	20769
SC/ST (=1)	0.214 [0.410]	0.185 [0.389]	-0.028	0.129	20695
OBC/EBC (=1)	0.597 [0.490]	0.612 [0.487]	0.016	0.511	20695
Smartphone (=1)	0.915 [0.279]	0.915 [0.279]	0.000	0.990	20890
Roof/terrace (=1)	0.540 [0.498]	0.525 [0.499]	-0.017	0.418	20811
Land owned (acres)	0.972 [1.597]	0.960 [1.502]	-0.007	0.900	20882
Flood-season income, 2022 (INR, thousands)	30.663 [40.496]	31.144 [42.484]	0.402	0.742	20569
Distance to nearest river (km)	2.754 [2.161]	2.643 [2.016]	-0.148	0.480	20890
Agricultural occupation (=1)	0.683 [0.465]	0.675 [0.468]	-0.006	0.717	20794
Agricultural occupation (on farm) (=1)	0.555 [0.497]	0.551 [0.497]	-0.002	0.906	20794
# Rooms in house	3.283 [1.939]	3.288 [1.967]	0.020	0.766	20402
Pakka house (=1)	0.408 [0.491]	0.400 [0.490]	-0.010	0.619	20399
Use toilet (=1)	0.722 [0.448]	0.723 [0.448]	-0.000	0.986	20869
# Years flooded (2017–2021)	1.772 [1.291]	1.862 [1.319]	0.091	0.387	20890
Time-varying characteristics (observed per round; averaged across rounds for full-sample columns)					
No flooding (=1)	0.761 [0.426]	0.777 [0.416]	0.017	0.524	20890
Low flooding (=1)	0.173 [0.378]	0.169 [0.375]	-0.005	0.818	20890
Severe flooding (=1)	0.066 [0.248]	0.054 [0.226]	-0.012	0.337	20890
Flood days (SAR)	1.590 [4.642]	1.369 [4.356]	-0.204	0.421	20890
F-stat (joint)			0.817	0.715	
Clusters	159	160			319

Notes: The unit of observation is a household-round in the pooled 2022–2025 sample. The Control and Treatment columns report unweighted means and standard deviations, in brackets, over non-missing household-round observations. All household characteristics are fixed at their 2022 value in every round. Flood categories and synthetic-aperture-radar (SAR) flood days are observed separately in each round. Variable definitions otherwise follow Table 1. The Difference column is the treatment coefficient from a separate OLS regression for each characteristic with district-round fixed effects. Reported p-values are based on standard errors clustered by gram panchayat, and N is variable specific. Among the time-varying rows, p-values below 0.10 are bolded. The joint F-statistic comes from a regression of treatment assignment on all listed characteristics with the same fixed effects and clustering. The sample contains 319 clusters.

Table 3. Alert receipt, timeliness, channels, content

Outcome	Treatment Effect	Control Mean	N	Clusters
Panel A. Alert receipt				
Any alert (=1)	0.093*** (0.011)	0.350	20,852	319
Total number of alerts	4.684*** (0.291)	2.115	20,790	319
Panel B. Channels				
In-person	-0.003 (0.007)	0.183	20,762	319
Outreach	0.016** (0.008)	0.089	20,762	319
IT-based	0.166*** (0.009)	0.128	20,762	319
Government	-0.008 (0.007)	0.172	20,762	319
Panel C. Timeliness, trust, and forecast error				
Time: before home (=1)	0.032*** (0.007)	0.111	20,087	319
Trust: completely (=1)	0.065*** (0.009)	0.208	20,194	319
Useful: very (=1)	0.062*** (0.010)	0.195	20,194	319
Accurate: very (=1)	0.066*** (0.010)	0.223	20,194	319
Alert, but no flood (=1)	0.035*** (0.007)	0.052	20,845	319
Flood, but no alert (=1)	-0.061*** (0.013)	0.439	20,845	319
Forecast error (=1)	-0.026** (0.010)	0.491	20,845	319
Shared information (=1)	0.106*** (0.016)	0.549	5,034	319
Shared information outside family (=1)	0.104*** (0.016)	0.536	5,034	319
Panel D. Alert content				
Date and time of flooding (=1)	0.071*** (0.006)	0.077	20,087	319
Forecasted water depth in the panchayat (=1)	0.051*** (0.006)	0.074	20,087	319
The depth of the river (=1)	0.036*** (0.005)	0.059	20,087	319
Level of flood severity (=1)	0.043*** (0.007)	0.126	20,087	319
River name (=1)	0.070*** (0.009)	0.238	20,087	319
That a flood was forecasted (=1)	0.060*** (0.009)	0.216	20,087	319
Increase/decrease in river water level (=1)	0.092*** (0.010)	0.292	20,087	319
Embankment broke (=1)	0.012* (0.006)	0.081	14,526	319
Water released from the dam (=1)	0.016 (0.010)	0.192	14,526	319

Notes: Each row reports the intent-to-treat coefficient from a separate OLS regression of the listed outcome on treatment assignment. Panel A is based on the questions "Did you receive any flood warning/alert?" and, for every month in which an alert was reported, "How many alerts did you receive that month?" Any alert equals one for an affirmative response, and Total number of alerts sums the month-specific counts. Panel B harmonizes responses to "What were the modes through which the flood alert was relayed to you in the month of [month]?" The response options are heard in person, loudspeakers, flags, phone call, WhatsApp/SMS, smartphone, computer, newspaper, television/radio, and public meeting; each broad channel indicator equals one if any constituent mode was reported. In Panel C, the timing indicator equals one when an alert was received before flood waters reached the respondent's home. The trust, usefulness, and accuracy indicators equal one for the top categories (completely, very useful, and very accurate, respectively) from the questions asking how much the respondent trusted the message and how useful and accurate the alert was. The forecast-error indicators classify an alert without a flood, a flood without an alert, and either mismatch. The same panel reports responses to "Did you discuss/share information regarding flooding with anyone else?" and whether the reported recipients discussed the alert with people in the same tola, other tola, other village, outside panchayat, or community leaders (classified as outside family). Panel D is based on the question asking what information the alert typically contained. Control means are unweighted means among control observations for the corresponding outcome. All regressions absorb district-round and midline-status fixed effects. Standard errors are clustered by gram panchayat and reported in parentheses; N and the number of clusters are outcome specific. Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 4. Protective action: indices, sub-indices, components

Panel A. Indices

	OLS			2SLS (IV)		
	(1) Adv. resp.	(2) Last-minute resp.	(3) All	(4) Adv. resp.	(5) Last-minute resp.	(6) All
Treatment (ITT)	0.052**	-0.009	0.042*			
	(0.022)	(0.026)	(0.023)			
Any alert = 1 (LATE)				0.570***	-0.085	0.461*
				(0.220)	(0.280)	(0.236)
Kleibergen-Paap F				70.75	70.64	70.75
District × Round FE	Yes	Yes	Yes	Yes	Yes	Yes
Flood-days FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	20771	20760	20771	20734	20723	20734

Panel B. Sub-indices

	(1) Health	(2) Animals	(3) House	(4) Agriculture	(5) Food/water
Treatment	0.028	0.039*	0.050***	0.031*	0.041**
	(0.024)	(0.023)	(0.019)	(0.017)	(0.020)
District × Round FE	Yes	Yes	Yes	Yes	Yes
Flood-days FE	Yes	Yes	Yes	Yes	Yes
Observations	15188	20760	20765	20766	20765

Panel C. Components

	(1) Drainage canal (=1)	(2) Sandbagging (=1)	(3) Stocked food (=1)
Treatment	0.010*	0.015*	0.023*
	(0.005)	(0.008)	(0.012)
Control mean (pooled)	0.081	0.191	0.426
Control mean 2023	0.132	0.179	0.341
Control mean 2024	0.054	0.240	0.494
Control mean 2025	0.066	0.153	0.439
District × Round FE	Yes	Yes	Yes
Flood-days FE	Yes	Yes	Yes
Observations	14706	15183	15183

Notes: The survey first asks, "Did your household undertake any measures to protect household members, household items, or livestock, or any other protective measures against flooding in [year]?" Respondents who report taking measures are then asked what steps they took. Panel A reports standardized advanced-response, last-minute-response, and overall protective-action indices. Columns (1)–(3) are intent-to-treat OLS estimates. Columns (4)–(6) are two-stage least-squares estimates in which treatment assignment instruments for receipt of any alert; the coefficient on alert receipt is the corresponding local average treatment effect under the usual IV assumptions. Kleibergen–Paap Wald rk F-statistics are reported below the IV estimates. Panel B reports OLS effects on standardized indices grouping actions related to health, animals, the house, agriculture, and food or water. Panel C reports three binary components: whether the household dug a drainage canal, secured the house or non-agricultural land with sandbags or other protective measures, and stocked food. Pooled and round-specific control means are calculated within each component's estimation sample; a round is shown whenever at least one of the three components is observed in that round. All regressions absorb district-by-round, exact SAR flood-days, and midline-status fixed effects. Standard errors are clustered by gram panchayat and reported in parentheses. Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 5. Household welfare, overall index

	Heterogeneous by flood severity	
	(1) Continuous	(2) Discrete
Treatment (ITT)	0.010	
	(0.028)	
Flood severity (# days)	-0.018***	
	(0.004)	
Low flooding		-0.183***
		(0.045)
Severe flooding		-0.359***
		(0.082)
Treatment × Flood severity	0.015***	
	(0.004)	
Treatment × No flood		0.001
		(0.028)
Treatment × Low flooding		0.073
		(0.056)
Treatment × Severe flooding		0.299***
		(0.100)
District × Round FE	Yes	Yes
p: per-day loss fully offset	0.308	
p: low flood fully offset		0.009
p: severe flood fully offset		0.425
p: treated outcomes, low = severe		0.468
Observations	20375	20375

Notes: The dependent variable in both columns is the standardized overall household-welfare index. The index combines the physical-health, asset, agricultural-harvest, food-security, and coping domains. Table 6 reports these subindices. Column (1) uses continuous SAR flood days. The Treatment coefficient is the treatment-control difference at zero flood days; Flood severity is the control-group association with one additional flood day; and Treatment × Flood severity is the additional treatment-control difference per flood day. The reported p-value for "per-day loss fully offset" tests whether the sum of the flood-days and interaction coefficients equals zero. Column (2) uses mutually exclusive no-, low-, and severe-flood categories. Low and severe flooding are measured relative to no flooding among control households, while the three treatment interactions are treatment-control differences within the corresponding flood category. The "low flood fully offset" and "severe flood fully offset" p-values test whether the sum of the relevant control-group flood-category coefficient and treatment interaction equals zero. The "treated outcomes, low = severe " p-value test whether the low and severe flooding offsets are equal to zero. Both regressions absorb district-by-round fixed effects. Standard errors are clustered by gram panchayat and reported in parentheses. Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 6. Household welfare: sub-indices (binned) and key outcomes

Panel A. Sub-indices

	(1) Health	(2) Asset	(3) Harvest	(4) Food security	(5) Coping
Low flooding	0.015	-0.175***	-0.302***	-0.122**	-0.064*
	(0.033)	(0.043)	(0.048)	(0.049)	(0.039)
Severe flooding	0.008	-0.370***	-0.481***	-0.103*	-0.174**
	(0.043)	(0.106)	(0.084)	(0.061)	(0.078)
Treatment × No flood	0.019	0.002	-0.008	-0.004	-0.008
	(0.021)	(0.030)	(0.039)	(0.030)	(0.022)
Treatment × Low flooding	-0.025	0.105*	0.134**	0.044	0.010
	(0.044)	(0.057)	(0.062)	(0.061)	(0.046)
Treatment × Severe flooding	0.053	0.217*	0.283***	0.158**	0.211**
	(0.063)	(0.119)	(0.100)	(0.074)	(0.089)
District × Round FE	Yes	Yes	Yes	Yes	Yes
Observations	20106	20375	14205	14524	20106

Panel B. Key outcomes

	(1) Home flooded (=1) (Index: asset)	(2) House damage (000s INR; index: asset)	(3) Food insecure (=1) (Index: food security)	(4) Loan amount (000s INR; index: coping)
Low flooding	0.137***	1.517***	0.035**	4.124
	(0.024)	(0.430)	(0.016)	(2.653)
Severe flooding	0.237***	2.347***	0.054**	7.762
	(0.048)	(0.526)	(0.026)	(4.837)
Treatment × No flood	0.009	-0.130	-0.007	2.129
	(0.015)	(0.268)	(0.010)	(1.604)
Treatment × Low flooding	-0.029	-0.000	-0.015	-1.908
	(0.035)	(0.588)	(0.021)	(3.181)
Treatment × Severe flooding	-0.127**	-1.907***	-0.067**	-12.632**
	(0.059)	(0.719)	(0.029)	(5.064)
Control mean (pooled)	0.147	2.956	0.205	36.345
District × Round FE	Yes	Yes	Yes	Yes
Observations	20363	19711	14524	20099

Notes: Each column reports a separate OLS regression using the discrete flood-severity specification in column (2) of Table 5. The low- and severe-flood coefficients compare control households in those categories with control households experiencing no flood. The treatment-by-category coefficients report the treatment-control difference within each flood category. Panel A reports standardized physical-health, asset, agricultural-harvest, food-security, and coping subindices. Panel B reports selected outcomes in levels. The 2025 home-flooding question asks, "Did flood waters enter your home between Aashadh and Kartik in 2025?"; earlier rounds use the corresponding year-specific reference periods. Home flooded equals one for a yes response. House damage is the response to "What was the total cost of damage or repairs to your house?" Food insecure equals one when, in response to the question whether lack of money or other resources created specified food concerns between Aashadh and Kartik, the respondent selects "You were hungry but did not eat." Loan amount is reported after the questions "Have you or your family taken any loans since Aashadh?" and "What amount of loan was taken during this period?". House damage and loan amount are expressed in thousands of Indian rupees. Pooled control means are calculated over control observations in the corresponding estimation sample. All regressions absorb district-by-round fixed effects. Standard errors are clustered by gram panchayat and reported in parentheses. Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 7. Government relief: indices, support types

Panel A. Relief indices

	(1) All relief	(2) Cash/food transfers	(3) Other supplies	(4) Evacuation/ rehabilitation
Low flooding	0.347***	0.371***	0.272***	0.240***
	(0.081)	(0.082)	(0.071)	(0.073)
Severe flooding	0.409***	0.390***	0.260***	0.396***
	(0.119)	(0.096)	(0.098)	(0.148)
Treatment × No flood	0.006	0.019	0.015	-0.019
	(0.052)	(0.052)	(0.042)	(0.045)
Treatment × Low flooding	-0.146	-0.131	-0.079	-0.164*
	(0.104)	(0.105)	(0.099)	(0.088)
Treatment × Severe flooding	-0.173	-0.053	-0.129	-0.260
	(0.139)	(0.121)	(0.113)	(0.163)
District × Round FE	Yes	Yes	Yes	Yes
Observations	10065	10065	10065	10065

Panel B. Support types

	(1) Ration (=1)	(2) Cash transfer (=1)	(3) Food from community kitchen (=1)
Low flooding	0.074***	0.104***	0.062***
	(0.018)	(0.025)	(0.018)
Severe flooding	0.070***	0.164***	0.115***
	(0.021)	(0.040)	(0.039)
Treatment × No flood	0.015	-0.011	-0.003
	(0.009)	(0.021)	(0.010)
Treatment × Low flooding	-0.018	-0.039	-0.056***
	(0.023)	(0.032)	(0.020)
Treatment × Severe flooding	-0.014	-0.040	-0.091**
	(0.029)	(0.052)	(0.041)
Control mean (pooled)	0.050	0.145	0.064
District × Round FE	Yes	Yes	Yes
Observations	10065	10065	10065

Notes: The survey asks, "Did you receive any support from the administration (such as the panchayat, block, district, state, or central government), politicians, NGOs, or other organizations during or after the [year] flood season?" Respondents answering yes are asked, "What kind of benefit did you receive?" Each column reports a separate OLS regression using mutually exclusive no-, low-, and severe-flood categories. The low- and severe-flood coefficients compare control households in those categories with control households experiencing no flood, and the treatment-by-category coefficients are treatment-control differences within each flood category. Panel A reports standardized indices. Cash/food transfers comprises dry rations, food packets, and cash transfers through direct benefit transfer (DBT); Other supplies comprises candles, electricity, and polythene sheets; and Evacuation/rehabilitation comprises food from a community kitchen, health camps, relief-shelter camps, and assisted evacuation or boats. The all-relief index aggregates these three domains. Each support item is set to zero when a household reports receiving no support. Items are standardized within round using the control-group mean and standard deviation. When at least one component is observed, missing components are imputed at the corresponding treatment-arm-by-round mean; component z-scores are then equally weighted, and each resulting index is re-standardized within round against the control group. Panel B reports binary indicators for selecting "Dry ration," "Cash transfer through DBT," and "Food from the community kitchen" from the benefit list. Pooled control means are calculated in the corresponding estimation sample. All regressions absorb district-by-round fixed effects. Standard errors are clustered by gram panchayat and reported in parentheses. Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 8. Beliefs about EWS, worry about flooding

	Continuous (avg severity)			Binned (years flooded)		
	(1) False alarm (=1)	(2) Missed alert (=1)	(3) Flood concern (=1)	(4) False alarm (=1)	(5) Missed alert (=1)	(6) Flood concern (=1)
Treatment (ITT)	0.003	-0.048***	-0.006			
	(0.015)	(0.017)	(0.017)			
Average flood severity (# days)	0.002	0.014***	0.010***			
	(0.003)	(0.004)	(0.002)			
Flooded 1 of last 4 years				0.013	0.051	0.096***
				(0.028)	(0.034)	(0.026)
Flooded 2 of last 4 years				-0.027	0.135***	0.059**
				(0.026)	(0.029)	(0.024)
Flooded 3–4 of last 4 years				-0.009	0.116***	0.118***
				(0.025)	(0.032)	(0.025)
Treatment × Flood severity	-0.008**	-0.018***	-0.009*			
	(0.004)	(0.005)	(0.005)			
Treatment × Never flooded				-0.004	-0.035	-0.000
				(0.018)	(0.021)	(0.021)
Treatment × Flooded 1 year				-0.007	-0.069*	-0.045
				(0.034)	(0.036)	(0.033)
Treatment × Flooded 2 years				0.010	-0.152***	-0.021
				(0.033)	(0.035)	(0.032)
Treatment × Flooded 3–4 years				-0.055*	-0.138***	-0.069**
				(0.030)	(0.038)	(0.035)
Control mean	0.275	0.317	0.783	0.275	0.317	0.783
District × Round FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	4990	5018	5034	4990	5018	5034

Notes: The table reports 2025 OLS estimates for three binary survey outcomes under two measures of flood history. False alarm equals one for a yes response to "In the past few years, did you receive any alerts that turned out to be false alarms (i.e., flood water did not arrive as predicted)?" Missed alert equals one for a yes response to "In the past few years, did you experience floods but did not receive an alert about it?" For flood concern, respondents are asked how worried they are about unexpected flooding in 2026 and choose among not at all concerned, slightly concerned, somewhat concerned, very concerned, and extremely concerned. Flood concern equals one for any positive level of concern and zero for "Not at all concerned". Columns (1)–(3) use average flood severity, defined as average flood days over the preceding four years. Flood severity here counts days of flooding of any intensity and is not restricted to villages classified as severely flooded. The Treatment coefficient is the treatment-control difference at zero average flood days; the severity coefficient is the control-group association with one additional average flood day; and the interaction reports how the treatment-control difference changes with average severity. Columns (4)–(6) use the number of years flooded over the same four-year window: never flooded, one year, two years, and three or four years. Never flooded is the omitted control-group category; the flood-frequency coefficients compare the other control groups with it, and the four treatment interactions are treatment-control differences within the corresponding frequency categories. Control means are calculated over control observations in each outcome's estimation sample. All regressions absorb district-by-round fixed effects. Standard errors are clustered by gram panchayat and reported in parentheses. Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 9. Science and technology

	Flooding changed views on science & technology a lot (=1)	Science & technology sentiment index	Strongly disagree: too much science, not enough faith (=1)	Strongly disagree: science erodes right vs. wrong (=1)	Strongly disagree: science unimportant in daily life (=1)	World better off from science & technology (=1)
Low flooding	-0.021	-0.031	-0.005	-0.013	-0.020	0.022
	(0.018)	(0.042)	(0.008)	(0.009)	(0.021)	(0.021)
Severe flooding	-0.026	0.048	-0.009	-0.018	0.047	0.062**
	(0.022)	(0.050)	(0.008)	(0.012)	(0.030)	(0.025)
Treatment × No flood	0.005	-0.026	-0.007	-0.003	-0.006	0.001
	(0.013)	(0.028)	(0.005)	(0.007)	(0.017)	(0.012)
Treatment × Low flooding	0.016	0.048	-0.008	0.017	0.034	0.007
	(0.024)	(0.050)	(0.009)	(0.012)	(0.025)	(0.025)
Treatment × Severe flooding	0.106***	0.117*	-0.013	0.033*	0.065**	0.028
	(0.028)	(0.061)	(0.010)	(0.017)	(0.031)	(0.028)
Control mean (pooled)	0.436		0.044	0.066	0.449	0.569
Control mean 2022	.		.	.	.	.
Control mean 2023	.		.	.	.	.
Control mean 2024	0.420		0.049	0.071	0.419	0.565
Control mean 2025	0.452		0.039	0.061	0.479	0.573
District × Round FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	10066	10066	10066	10066	10066	10066

Notes: The table reports separate OLS regressions for science-and-technology outcomes collected in 2024 and 2025. Low and severe flooding are measured relative to no flooding among control households; the three treatment-by-category coefficients report treatment-control differences within the no-, low-, and severe-flood categories. For column (1), the 2025 survey asks, "Have your perceptions of science and technology changed from the experience of flooding in 2025?" with response options a lot, fairly, little, nothing, and do not know; the 2024 version asks how the respondent's perceptions changed from the experience of that year's flooding using the same scale. The outcome equals one for "A lot." Column (2) is the standardized science-and-technology sentiment index used in the analysis. For columns (3)–(5), respondents are told that three statements will be read aloud and are asked how much they agree or disagree with each: "We depend too much on science and not enough on faith"; "One of the bad effects of science is that it breaks down people's ideas of right and wrong"; and "It is not important for me to know about science in my daily life." The response scale is strongly agree, agree, neither agree nor disagree, disagree, strongly disagree, and do not know; each reported outcome equals one for "Strongly disagree." Column (6) uses the question "How much better off or worse off do you think the world is because of science and technology?" The response options are a lot worse off, somewhat worse off, neither worse nor better off, somewhat better off, a lot better off, and do not know; the outcome equals one for "A lot better off." Pooled and year-specific control means are reported for the non-index outcomes within their estimation samples; dots denote rounds in which an outcome was not collected, and control means are intentionally omitted for the standardized index. All regressions absorb district-by-round fixed effects. Standard errors are clustered by gram panchayat and reported in parentheses. Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 10. Demand for flood delivery service

	OLS (ITT)				2SLS (LATE)	
	(1) Practice WTP	(2) Monsoon WTP	(3) Flood WTP	(4) Flood WTP	(5) Flood WTP	(6) Flood WTP
Treatment (ITT)	-0.111	0.031	1.809*	1.739**		
	(0.621)	(1.096)	(1.088)	(0.757)		
Any alert = 1 (LATE)					12.886	12.539**
					(7.916)	(5.604)
Practice-round WTP		0.886***	1.090***		1.075***	
		(0.026)	(0.025)		(0.026)	
Monsoon WTP				0.760***		0.756***
				(0.010)		(0.010)
Control mean (unconditional)	28.06	28.74	33.94	33.94	33.92	33.92
Control mean (WTP >0)	35.34	63.70	64.69	64.69	64.70	64.70
% control with WTP >0	79.4	45.1	52.5	52.5	52.4	52.4
Kleibergen-Paap F					66.45	66.10
District $\times$ Round FE	Yes	Yes	Yes	Yes	Yes	Yes
Avg. Flood-Days FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	5034	5034	5034	5034	5029	5029

Notes: The 2025 survey elicits maximum willingness to pay (WTP) in Indian rupees using a random-price procedure. Respondents are told to state the highest price they are willing and able to pay; then a price is drawn independently of the stated bid. If the draw is less than or equal to the bid, the respondent can purchase at the drawn price, and payment for an actual service is collected in cash. A hypothetical practice round asks the maximum amount the respondent is willing to pay for a notebook with a market price of Rs 50. The Flood Alert Service provides WhatsApp or SMS messages from June through October 2026 with the forecasted water depth, expected timing of flooding in the village, and latest water level of the nearest river, generally two to four days before a potential flood. The Monsoon Forecast Service provides farmer-focused WhatsApp or SMS forecasts derived from Indian Meteorological Department forecasts, including the predicted window for sustained monsoon onset, how early or late that is relative to a normal year, forecast confidence, and updates on unusual pauses or restarts; messages may begin about 30 days before onset. Each service has a stated market-price ceiling of Rs 100. The final elicitation for each product asks, "What is the maximum amount you are willing to pay to buy [the notebook/this flood alert service/this Monsoon Forecast Service]?" Column (1) reports the intent-to-treat effect on practice-round WTP. Column (2) reports the effect on monsoon-service WTP while controlling for practice-round WTP. Columns (3) and (4) report effects on flood-alert-service WTP while controlling for practice-round and monsoon-service WTP, respectively. Columns (5) and (6) are the corresponding two-stage least-squares versions of columns (3) and (4), with receipt of any alert instrumented by treatment assignment; Kleibergen–Paap Wald rk F-statistics are reported below. Control means are shown both unconditionally and among positive-WTP observations, together with the control share reporting positive WTP. All regressions absorb district-by-round and average historical flood-days fixed effects. Standard errors are clustered by gram panchayat and reported in parentheses. Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 11. Cost-effectiveness

Panel A. Benefits (single channel)	
Averted distress borrowing per household under severe flood (Rs)	12632
× P(severe flood), 2022–25 realized	5.4%
× households covered (160 GPs × 2000)	320000
× flood seasons	4
Total benefits	**Rs 873.1M**
Panel B. Costs (measured, delivery only)	
Cost per GP per season (Rs)	46509
× GP-seasons (160 × 4)	640
Total cost	**Rs 29.8M**
per household-season	Rs 23
Benefit–cost ratio	**29:1**

Notes: Panel A monetizes a single measured benefit channel: Rs 12,632 in averted borrowing per household under severe flooding, based on the treatment-by-severe-flood estimate for loan amount in Table 6. The calculation multiplies this amount by the 5.4 percent realized probability of severe flooding in 2022–2025, 320,000 households in 160 treated gram panchayats, and four flood seasons. Panel B calculates measured delivery costs as Rs 46,509 per gram-panchayat season, covering agent stipends and NGO overhead, multiplied by 640 gram-panchayat seasons. Total benefits of Rs 873.1 million divided by total costs of Rs 29.8 million yield the reported benefit-cost ratio of 29:1 after rounding; measured cost per household-season is Rs 23. The calculation does not monetize other estimated outcome changes, including asset damage, harvests, food security, and public relief. Applying the 15.7 percent severe-flood incidence observed over 2017–2025 raises the benefit-cost ratio to approximately 85:1. The expected state relief-transfer benchmark is Rs 7,000 multiplied by 15.7 percent. The Rs 34 willingness-to-pay benchmark is the control-group mean maximum bid for the 2026 Flood Alert Service elicited with the random-price procedure described in the notes to Table 10.

11 Appendix

11.1 Appendix: Data

11.1.1 Survey Rounds and the 2022 Midline Sample

Within each round, the intervention runs for five months, June to October, coinciding with Bihar's monsoon season. The midline phone survey window is August 1 to September 30; the endline is conducted in person after the flood season ends. In the 2024 and 2025 rounds, the household surveys are complemented by a GP-level phone survey of mukhiyas (elected heads of the gram panchayat) and ward members (elected representatives of individual wards).

Because the household panel was defined at the 2022 endline, the first midline, conducted during the 2022 flood season, drew on a separate sample: we contacted mukhiyas and ward members in each treatment and control GP and asked them to provide phone numbers of resident households, whom we then surveyed. The 2022 midline sample is therefore a panchayat-leader-referred sample of residents rather than the tracked panel; we include it in the analysis of ex-ante outcomes. From 2023 onward, the midline was administered to the tracked panel, and households that did not complete a midline answered the relevant questions retrospectively at the subsequent in-person endline.

11.1.2 Sampling Protocol and Sample Characteristics

Within each of two randomly selected villages per GP, enumerators began at the village center and approached every third household until ten households were enrolled, subject to the smartphone and SC/ST quotas described in the main text. Where a quota would otherwise have gone unmet, enumerators continued the walk, replacing the sampled household with the next eligible one. Surveys were administered in person, in Hindi, using computer-assisted personal interviewing; midlines were conducted by phone.

Beyond the characteristics summarized in the main text, the average household includes 2.4 children under 18 and is headed by a 45-year-old with 8.8 years of schooling; 55% work on their own farm. Homes occupy 3.3 rooms on average, and 54% have access to a roof or terrace, the usual refuge for moving people and possessions above floodwaters.

11.1.3 Satellite Flood Measurement: Sensor, Imputation, and Validation

Sensor. The primary signal comes from Sentinel-1, available from 2017, a satellite constellation operated by the European Space Agency carrying C-band SAR. SAR is an active sensor: it transmits microwave pulses and measures the energy scattered back. Smooth open water acts as a specular reflector, deflecting the pulse away from the sensor and returning a distinctively low backscatter signal, which makes standing water sharply identifiable. Because the radar supplies its own illumination and microwaves pass through cloud, Sentinel-1 observes the surface at night and through the near-continuous cloud cover of the monsoon, precisely the conditions under which optical satellites are blind. For each pass over a study village, we classify pixels as water or non-water and record the fraction of the village's pixels inundated.[21]

[21] For 35 of the 591 sampled villages, village boundaries could not be reliably matched in the administrative shapefile. For these villages we assign, on each day, the average inundation of measured villages in the same GP, typically

Imputation. Following Patel (2024), we fill the ∼81% of village-days without a direct SAR pass using a Random Forest model trained on directly observed days. For each unobserved village-day, the model predicts the fraction of pixels flooded from predictors observed daily or near-daily: lagged SAR backscatter, by far the most informative input (62% of feature importance); soil moisture; river gauge readings from the CWC stations described in the main text; and optical water indices from Landsat, Sentinel-2, and PlanetScope imagery on cloud-free days. The output is a continuous daily panel of village-level inundation spanning all nine monsoon seasons.

Validation. On a held-out set of 119,010 village-days with direct SAR observations, the imputation achieves an R^2 of 0.93 with a mean absolute error of 1.8 percentage points of village area. Performance is hardest during the monsoon, when inundation is most variable, yet remains strong: $R^2 = 0.92$ and a mean absolute error of 3.2 percentage points in monsoon months, against 0.95 and 0.8 percentage points outside them.

Threshold and episode measures. The 30-percentage-point flood-day threshold is deliberately demanding on two grounds. It is roughly an order of magnitude larger than the imputation model's typical monsoon error, so a flood day cannot plausibly be an artifact of measurement noise. And it is physically unambiguous: nearly a third of a village newly under water, relative to its own seasonal norm, is flooding that reaches homesteads and fields, not marginal waterlogging or classification jitter at the edges of ponds and rivers. Because a single flood can register as several bursts of flood days separated by brief recessions, we also construct episode-level measures: grouping flood days separated by at most seven dry days into single events, we record each village-year's number of flood events and their average and longest durations. Total episode duration and the raw flood-day count correlate at 0.98 across village-years, and results are robust to alternative thresholds (20 and 40 percentage points) and to using bridged episode duration in place of the flood-day count. Over the full 2017–2025 panel, flooding is more common than in the study window: 37.5% of village-years record at least one flood day and 16% are severe, reflecting the floodier 2017–2021 pre-period.

11.1.4 Flood Realizations: Season Detail and Village Histories

Figures 3(b) and (c) summarize incidence and duration by season, and Table A1 reports village histories. Among the 591 villages observed in all four seasons, 19% flooded in 2022 (3% severely), 13% in 2023 (one village severely, averaging 2.2 flood days conditional on flooding), 29% in 2024 (7% severely), and 33% in 2025 (15% severely, averaging 10.7 flood days). Pooled, 76.6% of village-years experience no anomalous flooding, 17.2% low, and 6.2% severe.

Village histories are varied: 271 villages flood at least once, 185 more than once, and only 17 in every season; 104 flood severely at least once, 34 more than once, and none in all four. The modal severe history is a single hit in 2025 (Table A1, Panel B).

The pre-study peak came in 2020, when 37% of villages were severely flooded and flooded villages averaged 14.3 flood days; roughly half of villages flooded in a typical pre-study year, against under a quarter during the study.

the co-sampled village (18 villages); where no measured village exists in the GP, the block average (14 villages); and failing that, the district average (3 villages).

Figure 3(a) plots within-season timing. In the seasons supplying most severe exposure, flooding concentrated in August and September: 2025 peaked in mid-August, when 11% of sample villages were simultaneously inundated, and 2024 brought successive waves through late August and September. The midline window (August 1–September 30) sits squarely inside this period, so household actions and alert receipt are measured while the season's floods are unfolding.

11.1.5 Outcome Components and Question Wording

Alert delivery and the information environment. From the midline we measure, first, delivery: whether the household received any flood alert during the season, the number of alerts received, whether alerts were timely (received before floodwaters reached their home), the mode of receipt (in-person, government media, community outreach, or IT-based channels), and the informational content conveyed (for example, the forecast date and time of flooding, expected water depth, and severity level). Second, quality: crossing each household's alert receipt against the household's own record of whether its GP in fact flooded yields two forecast-error indicators , a *false alarm* (alert received, no flood) and a *missed flood* (flood, no alert), the latter being the dangerous error an EWS exists to prevent. Third, credibility: households rate their trust in the alerts they received and the alerts' usefulness ("Typically, how much do you trust / how useful was the message you received across the months of Aashadh to Kartik?"), from which we construct indicators for the top response categories, trusting alerts *completely* and finding them *very* useful. We also capture whether households shared flood information with people in their neighbourhood.

Ex-ante adaptation. The midline asks: "Did your household undertake any measures to protect household members, household items, or livestock, or any other protective measures between Jeshth to Kartik? What steps did you take?", with responses recorded against a structured list of protective steps. We partition these steps into two meta-domains whose distinction carries the paper's central behavioral hypothesis. *Advance responses* are preparatory steps taken well ahead of flood onset, twenty-seven steps across five sub-domains: human health (six steps, e.g., stocking medicines and sanitary equipment), animal-related (two, e.g., vaccinating livestock), house and environment (six, e.g., digging a drainage canal, sandbagging), agricultural (nine, e.g., moving crops to higher ground, harvesting early), and food and water (four, e.g., stocking food, storing drinking water). *Last-minute responses* are rapid actions at or during flood onset, **five** steps such as moving household members to a safe place and evacuating documents and valuables. If warnings work by relaxing an information constraint, they should shift action from the second category into the first, increasing advance preparation without increasing scrambling at onset.

Ex-post welfare. From the endline we construct five domains of realized welfare. *Physical health* (**seven** components) includes the share of household members who fell sick and spending on treating illness. *Asset damage* (**seven**) includes whether the house flooded, the cost of house damage, and animal sickness. *Agricultural harvest* (**two**) measures harvest quantities in kilograms of the region's staple, paddy, and whether output for any crop grown was negatively affected relative to a no-flood year. *Food security* (**eight**) includes skipped meals and running out of food.

Coping (**seven**) includes postponing large purchases, drawing down savings, and borrowing, with loan amounts in rupees allowing a money-metric reading of distress finance. We additionally record receipt of government relief, dry rations, the state's INR 7,000 post-flood cash transfer, and meals from community kitchens, both as outcomes in their own right and to interpret welfare effects net of the state's response.

Beliefs and valuation. The 2024 and 2025 endlines measure whether the season's flooding changed the household's views of science and technology (an indicator for "a lot") and a four-component sentiment index adapted from the World Values Survey science module: whether the world is better off because of science and technology, and strong disagreement with the statements that science is unimportant in daily life, that we depend too much on science and not enough on faith, and that science erodes ideas of right and wrong. The 2025 endline adds an incentivized Becker–DeGroot–Marschak elicitation of willingness to pay for a flood-alert delivery service, benchmarked against a non-incentivized practice round and an incentivized monsoon-onset forecast (Dizon-Ross and Jayachandran, 2022).

11.2 Attrition Detail

Considering any contact (completion of the midline or the endline) we reach 88.1% of baseline households in 2023, 93.7% in 2024, and 92.4% in 2025 (Table A3, Panel A). Endline completion specifically is 79.9%, 90.1%, and 90.2% (Panel B); the 2023 round is the low-water mark on both margins, with completion rising by roughly ten percentage points as re-contact protocols matured. Regressing a completion indicator on treatment with district-by-round fixed effects, clustered at the GP, the treatment–control difference is insignificant in every round on every margin: p-values range from 0.27 to 0.65 for any contact and from 0.27 to 0.58 for endline completion, and aggregate tests for completion of all three follow-up rounds likewise do not reject ($p = 0.29$ for any contact, $p = 0.45$ for the endline).

11.3 Robustness of the severity measure

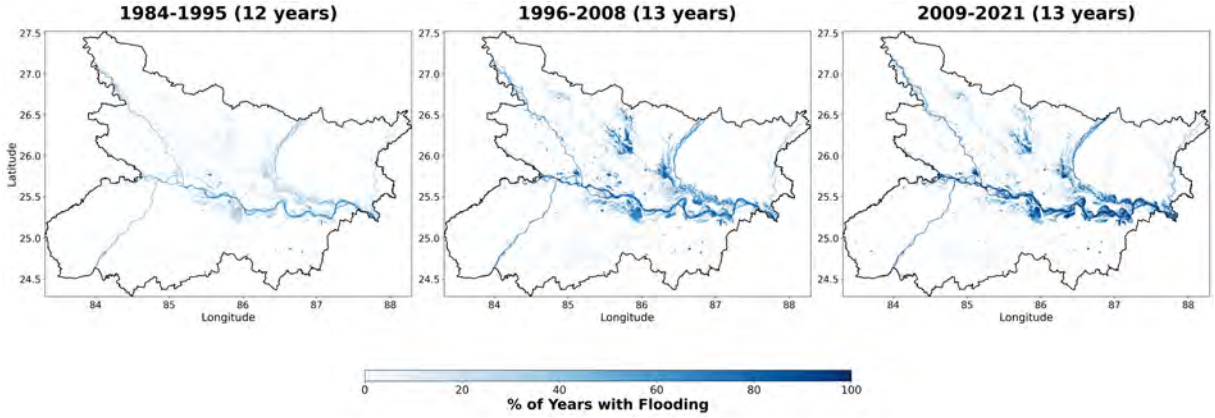
Our severity measure requires one discretionary input: the inundation threshold above which a village-day counts as flooded. Table A4 re-estimates the continuous specification of equation (2) for the overall welfare index at ten thresholds, from 0.05 to 0.50 in steps of 0.05, where a day adds to S_{vt} if the share of village pixels inundated, net of the village's seasonal mean, exceeds the threshold; column (6), at 0.30, is the definition used throughout the paper. The dose-response pattern is stable: the severity gradient is negative and significant at the 1 percent level at every threshold, and the treatment interaction is positive at every threshold, significant at the 1 percent level throughout the interior band from 0.20 to 0.35. Both coefficients grow in magnitude as the threshold rises because a flood-day then denotes more severe inundation, a relabeling of units rather than a change in the result. The two endpoints show why the interior band is where the measure belongs: at 0.05, 77 percent of village-seasons register positive severity, so a flood-day is a mild event and the per-day coefficients are small, while at 0.50 fewer than 8 percent do and precision collapses, with the interaction standard error rising monotonically from 0.001 to 0.012 across the columns. The implied share of the per-day welfare loss offset by treatment rises from under one-third at the lowest threshold to more than four-fifths at the paper's

definition, the continuous counterpart of the binned finding that protection is more complete under severe flooding. Finally, the uninteracted treatment coefficient rises with the threshold, from -0.009 to 0.026, while remaining insignificant throughout; this is a composition effect, since as the threshold rises the zero-severity cell absorbs mildly flooded village-years where treatment has real effects, so that coefficient no longer isolates unflooded village-years.

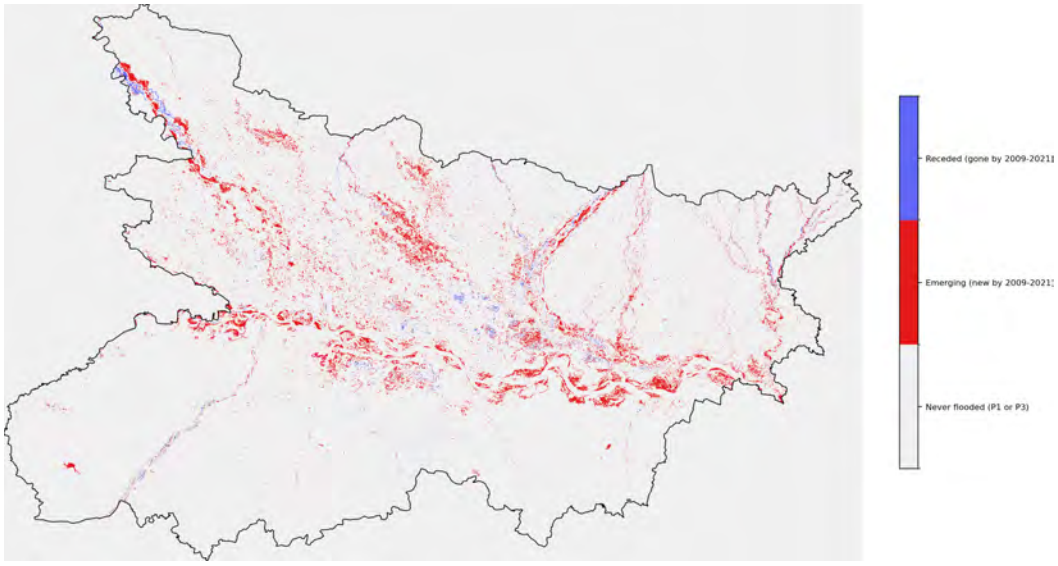
The measure also involves a second construction choice, how to treat brief gaps in the satellite record. Column (11) replaces the flood-day count with the total duration in days of bridged flood episodes, in which short gaps between flooded days are absorbed into a single episode. Both the severity gradient (-0.007, s.e. 0.002) and the treatment interaction (0.006, s.e. 0.002) remain significant at the 1 percent level, and treatment continues to offset a share of the per-day loss similar to the headline. The coefficients are smaller in absolute value than in column (6) because a day of episode duration, which includes within-episode gaps, is a milder unit than a threshold-qualifying flood-day.

11.4 Appendix: Figures and Tables

Appendix Figure A1. Historical flooding in Bihar, 1984–2021 (JRC)



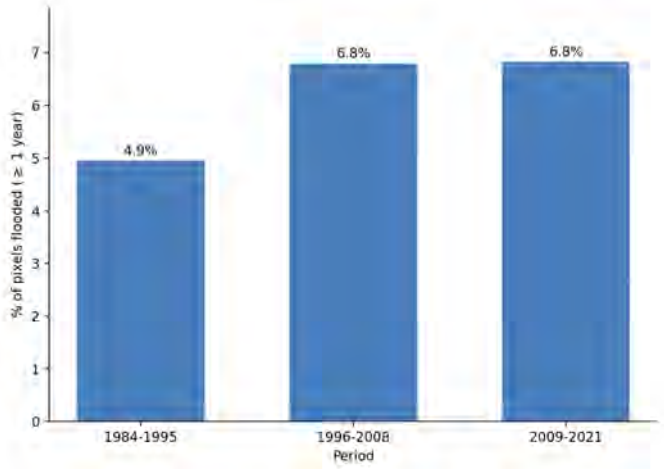
(a) Historical flood extent, 1984–2021



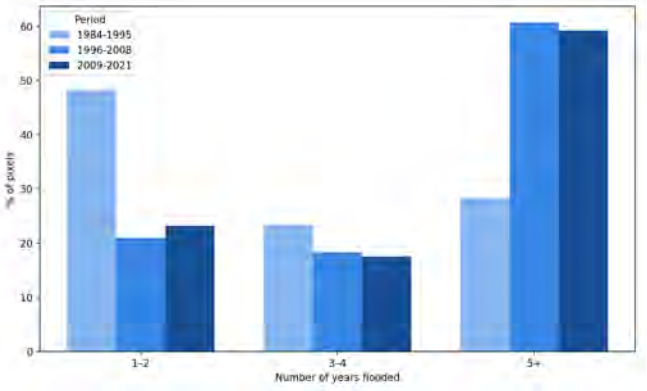
(b) Change in flood extent, 1984–2021

Notes: The figure uses Joint Research Centre (JRC) historical surface-water rasters for Bihar. A pixel is classified as flooded in a given year when the raster records seasonal or permanent surface water. Panel (a) divides 1984–2021 into three periods (1984–1995, 1996–2008, and 2009–2021) and maps the percentage of years in each period in which each pixel was flooded. Panel (b) compares flood status in the first and third periods. Emerging pixels were not flooded during 1984–1995 but were flooded at least once during 2009–2021; receded pixels were flooded at least once during 1984–1995 but were not flooded during 2009–2021. The remaining category comprises pixels whose flood status did not change between the two periods, including pixels flooded in both periods and pixels flooded in neither. The transition classification is restricted to pixels with valid data in all three periods.

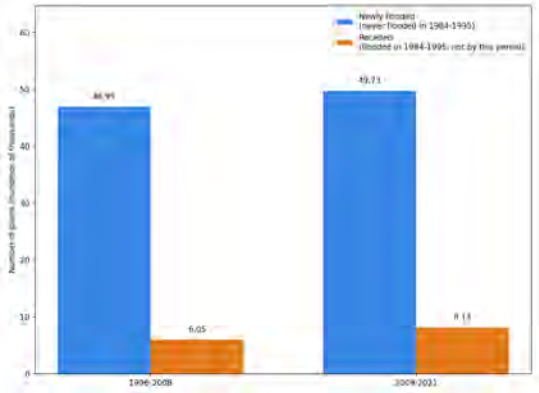
Appendix Figure A2. Flood frequency and transitions in Bihar, 1984–2021 (JRC)



(a) Share of pixels flooded at least once



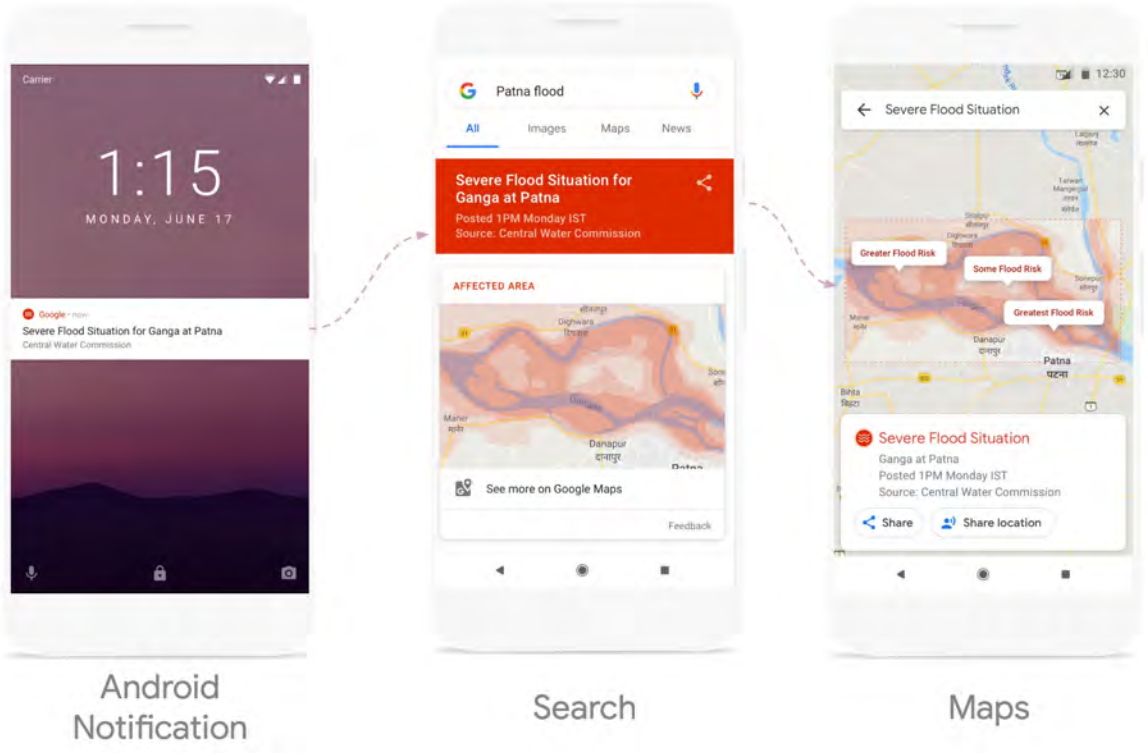
(b) Flood intensity, conditional on flooding



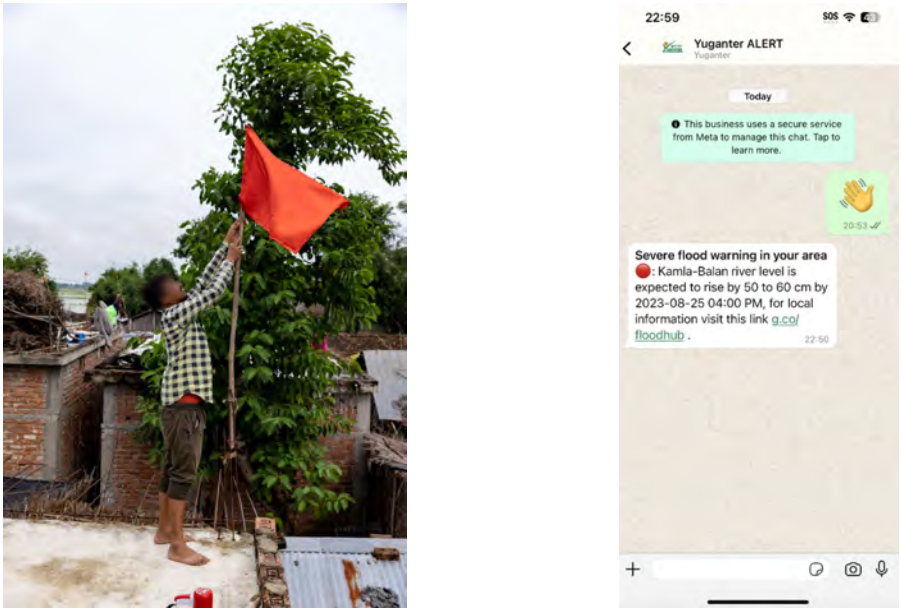
(c) Flooded conditional on no flooding, 1984–95

Notes: The figure uses Joint Research Centre (JRC) historical surface-water rasters for Bihar and divides 1984–2021 into three periods: 1984–1995 (12 years), 1996–2008 (13 years), and 2009–2021 (13 years). A pixel is classified as flooded in a given year when the raster records seasonal or permanent surface water, and each period raster stores the number of years in which a pixel was flooded. Panel (a) reports the percentage of pixels with valid data in a period that were flooded in at least one year of that period. Panel (b) restricts to pixels flooded at least once and reports the distribution of the number of years flooded, grouped as 1–2, 3–4, and 5 or more; shares are normalized within period, so the three bars of a period sum to 100 percent. Panel (c) reports pixel counts, in hundreds of thousands, rather than shares. Newly flooded counts pixels that were never flooded during 1984–1995 and were flooded at least once in the period shown; Receded counts pixels that were flooded at least once during 1984–1995 and were not flooded in the period shown. Both series are defined relative to 1984–1995, so only 1996–2008 and 2009–2021 appear. Panels (a) and (b) use the pixels with valid data in each period separately, while panel (c) retains only pixels with valid data in all three periods.

Appendix Figure A3. The intervention



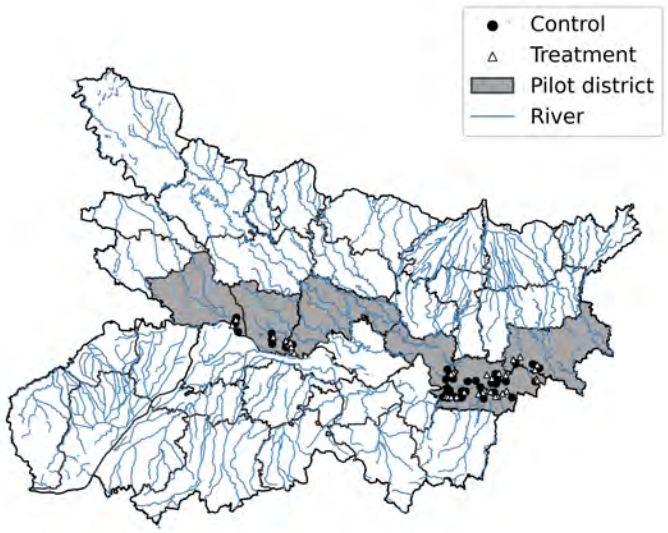
(a) Google notification: Android, Search, Maps



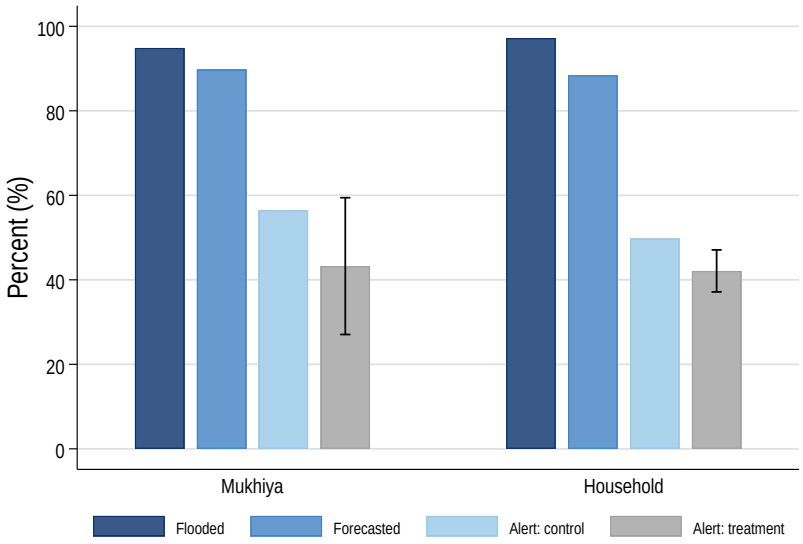
(b) Community agents: flag and WhatsApp group

Notes: Panel (a) displays examples of the three Google modes of flood-information dissemination: an Android flood notification, a Google Search flood-forecast card, and a Google Maps flood view. Panel (b) displays two community-level dissemination mechanisms used in treatment gram panchayats: flags and agent-run WhatsApp groups.

Appendix Figure A4. 2019 study: forwarding alerts to local leaders



(a) Pilot area: rivers, pilot districts, 77 GPs by arm



(b) Flooding, forecast coverage, and alert receipt

Notes: The figure summarizes a randomized intervention conducted during the 2019 monsoon, in which Google flood alerts were forwarded to mukhiyas by WhatsApp and SMS in half of 77 pilot gram panchayats across six districts. Panel (a) maps the pilot area: black circles and open triangles mark pilot gram panchayats assigned to the control and treatment arms, respectively; blue lines show rivers; and gray shading identifies the pilot districts. Panel (b) uses the January 2020 post-season survey and reports means for two groups of respondents: mukhiyas, and all surveyed respondents. Flooded and Forecasted are calculated among control respondents only. Flooded is the share reporting flooding during the 2019 season, and Forecasted is the share whose village appears in the pilot alert-generation lists, so it measures forecast coverage rather than anything the respondent reported. Alert: control and Alert: treatment are the shares of respondents in each arm reporting receipt of any flood alert; respondents with no recorded answer to the alert-receipt question are coded as not having received one. Whiskers are 95 percent confidence intervals.

Appendix Figure A5. Validation: all stations



Notes: The figure extends the validation exercise in Figure 5 to all 25 sample stations in 2025. Within each station panel, the black series, read on the left axis, is the percentage of mapped sample villages classified as flooded by SAR; open circles mark every tenth daily observation. The solid blue and dashed orange series, read on the right axis, are realized and forecasted water levels above the station-specific severe threshold. Solid red vertical lines with triangles mark the start of Google severe events, and dashed green vertical lines with squares mark the start of FloodHub severe events.

Appendix Table A1. Village flood histories, 2022–2025

2022	2023	2024	2025	N	%
Panel A. Any flooding					
NF	NF	NF	NF	320	54.1
NF	NF	F	F	60	10.2
F	NF	F	F	48	8.1
NF	NF	NF	F	44	7.4
NF	F	F	NF	28	4.7
F	NF	NF	NF	17	2.9
F	F	F	F	17	2.9
NF	F	NF	NF	13	2.2
F	F	F	NF	12	2.0
NF	NF	F	NF	12	2.0
F	NF	NF	F	8	1.4
F	NF	F	NF	4	0.7
F	F	NF	NF	4	0.7
NF	F	F	F	3	0.5
NF	F	NF	F	1	0.2
F	F	NF	F	0	0.0
Total				591	100.0
Panel B. Severe flooding					
NF	NF	NF	NF	487	82.4
NF	NF	NF	F	53	9.0
NF	NF	F	F	24	4.1
NF	NF	F	NF	9	1.5
F	NF	F	F	9	1.5
F	NF	NF	NF	7	1.2
F	NF	NF	F	1	0.2
NF	F	NF	NF	1	0.2
F	F	F	NF	0	0.0
F	F	F	F	0	0.0
F	NF	F	NF	0	0.0
F	F	NF	NF	0	0.0
NF	F	F	F	0	0.0
NF	F	F	NF	0	0.0
F	F	NF	F	0	0.0
NF	F	NF	F	0	0.0
Total				591	100.0

Notes: Each row is a four-season sequence of village flood status, where F denotes flooded and NF denotes not flooded. Panel A classifies a village-year as flooded when synthetic-aperture-radar (SAR) flood days exceed zero. Panel B classifies a village-year as severely flooded when the severe-flood indicator equals one. N is the number of villages with the indicated sequence, and % is that sequence's percentage of villages in the panel. All 16 possible sequences are displayed, including sequences containing zero villages; within each panel, sequences are ordered from most to least frequent.

Appendix Table A2. Balance, endline sample

	Control Mean [SD]	Treatment Mean [SD]	Difference	p-value	N
Time-invariant characteristics (imputed from 2022 baseline)					
# Males	2.816 [1.410]	2.805 [1.480]	-0.012	0.811	20106
# Females	2.681 [1.531]	2.652 [1.587]	-0.027	0.556	20106
# Under 18s	2.351 [1.852]	2.345 [1.930]	-0.003	0.966	20106
HoH age	46.051 [15.166]	46.177 [15.516]	0.105	0.824	20106
Mean age	30.341 [12.041]	30.894 [12.549]	0.529	0.226	20106
HH member 65+ (=1)	0.447 [0.677]	0.478 [0.701]	0.030	0.124	20106
HoH years of education	8.776 [3.500]	8.803 [3.433]	0.018	0.894	20102
Muslim (=1)	0.093 [0.291]	0.072 [0.258]	-0.020	0.262	19990
SC/ST (=1)	0.213 [0.410]	0.186 [0.389]	-0.027	0.140	19943
OBC/EBC (=1)	0.598 [0.490]	0.614 [0.487]	0.017	0.498	19943
Smartphone (=1)	0.914 [0.280]	0.915 [0.280]	0.000	0.947	20106
Roof/terrace (=1)	0.572 [0.495]	0.558 [0.497]	-0.015	0.445	20066
Land owned (acres)	0.962 [1.540]	0.975 [1.525]	0.018	0.737	20101
Flood-season income, 2022 (INR, thousands)	30.547 [40.421]	31.076 [42.347]	0.467	0.702	19801
Distance to nearest river (km)	2.760 [2.165]	2.651 [2.022]	-0.145	0.492	20106
Agricultural occupation (=1)	0.725 [0.447]	0.726 [0.446]	0.002	0.877	20080
Agricultural occupation (on farm) (=1)	0.537 [0.499]	0.530 [0.499]	-0.005	0.786	20080
# Rooms in house	3.358 [1.986]	3.370 [1.973]	0.028	0.629	19868
Pakka house (=1)	0.465 [0.499]	0.449 [0.497]	-0.018	0.344	19867
Use toilet (=1)	0.728 [0.445]	0.728 [0.445]	-0.001	0.949	20096
# Years flooded (2017–2021)	1.770 [1.290]	1.867 [1.316]	0.096	0.360	20106
Time-varying flood exposure (observed per round; averaged across rounds for full-sample columns)					
No flooding (=1)	0.760 [0.427]	0.776 [0.417]	0.017	0.531	20106
Low flooding (=1)	0.174 [0.379]	0.169 [0.375]	-0.006	0.800	20106
Severe flooding (=1)	0.066 [0.249]	0.055 [0.227]	-0.011	0.374	20106
Flood days (SAR)	1.594 [4.620]	1.386 [4.390]	-0.188	0.464	20106
F-stat (joint)				0.714	
Clusters	159	160			319

Notes: The unit of observation is a household-round, and the sample consists of observations coded as having completed the endline in the pooled 2022–2025 data. Control and Treatment report unweighted means and standard deviations, in brackets, over non-missing observations for each characteristic. All household characteristics are fixed at their 2022 value in every round. Flood categories and synthetic-aperture-radar (SAR) flood days are observed separately in each round. Variable definitions follow Table 1. Difference is the coefficient on treatment from a separate OLS regression for each characteristic with district and round fixed effects; the corresponding p-value uses standard errors clustered by gram panchayat, and N is variable-specific. The joint F-statistic comes from a complete-case regression of treatment assignment on all listed characteristics with the same fixed effects and clustering. The Clusters row reports the number of gram panchayats represented in the control group, treatment group, and full sample.

Appendix Table A3. Attrition

	Control Completed (%)	Treatment Completed (%)	Total Completed (%)	*p*-value (balance)
Panel A. Any follow-up completed				
Round 2023	88.6	87.6	88.1	0.296
Round 2024	93.5	93.9	93.7	0.654
Round 2025	93.1	91.7	92.4	0.104
All rounds (aggregate)	83.1	80.5	81.8	0.286
Panel B. Endline completed				
Round 2023	80.5	79.2	79.9	0.266
Round 2024	89.9	90.4	90.1	0.577
Round 2025	90.7	89.7	90.2	0.304
All rounds (aggregate)	73.9	72.0	72.9	0.445
Observations	2744	2838	5582	

Notes: The base sample consists of households observed in the 2022 baseline. Panel A reports the percentage of baseline households present in the analysis data in either the midline or endline follow-up round; Panel B reports the percentage coded as having completed the endline module in each follow-up round. The 2023, 2024, and 2025 rows are round-specific. In each panel, All rounds (aggregate) is the percentage of baseline households completing either the midline or endline (Panel A) or completing the endline module (Panel B) in all three follow-up rounds. For an individual-round row, the balance p-value is from an OLS regression of that round's indicator on treatment with district fixed effects. For the aggregate row, the reported p-value is instead from a stacked household-round regression of the round-specific indicator on treatment with district-by-round fixed effects. All standard errors are clustered by gram panchayat. Observations reports the number of baseline households in the control group, treatment group, and full sample.

Appendix Table A4. Flood measure robustness

	Anomaly threshold (share of village area)										Bridged episodes
	(1) 0.05	(2) 0.10	(3) 0.15	(4) 0.20	(5) 0.25	(6) 0.30	(7) 0.35	(8) 0.40	(9) 0.45	(10) 0.50	(11) Episode duration
Treatment (ITT)	-0.009	-0.002	0.004	0.005	0.006	0.010	0.015	0.021	0.023	0.026	0.013
	(0.029)	(0.028)	(0.028)	(0.028)	(0.028)	(0.028)	(0.028)	(0.029)	(0.028)	(0.028)	(0.028)
Flood severity (# days)	-0.007***	-0.009***	-0.010***	-0.012***	-0.016***	-0.018***	-0.021***	-0.022***	-0.026***	-0.029***	-0.007***
	(0.001)	(0.001)	(0.002)	(0.002)	(0.003)	(0.004)	(0.004)	(0.005)	(0.006)	(0.008)	(0.002)
Treatment × Flood severity	0.002*	0.004**	0.005**	0.008***	0.012***	0.015***	0.017***	0.017**	0.018*	0.021*	0.006***
	(0.001)	(0.002)	(0.002)	(0.003)	(0.004)	(0.004)	(0.006)	(0.007)	(0.010)	(0.012)	(0.002)
District × Round FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Village-seasons with severity > 0 (%)	76.9	57.6	45.2	35.5	29.0	23.4	18.1	14.4	10.6	7.8	23.4
Observations	20375	20375	20375	20375	20375	20375	20375	20375	20375	20375	20375

Notes: The dependent variable in every column is the standardized overall household-welfare index, which summarizes the physical-health, coping, asset-damage, food-security, and agricultural-harvest domains reported in Table 6. Each column reproduces the continuous flood-severity specification in Table 5 using an alternative synthetic-aperture-radar (SAR) flood measure. Columns (1)–(10) measure flood severity as the number of days on which the share of village area exhibiting a flood anomaly exceeds the indicated threshold; the thresholds run from 0.05 to 0.50 in increments of 0.05. Column (11) measures flood severity as the total duration, in days, of bridged flood episodes, with missing episode-duration values coded as zero. In every column, Treatment (ITT) is the treatment-control difference when measured flood severity equals zero; Flood severity is the control-group association between an additional measured flood day and household welfare; and Treatment × Flood severity is the corresponding change in the treatment-control difference per additional day. All regressions absorb district-by-round fixed effects. Standard errors are clustered by gram panchayat and reported in parentheses. Village-seasons with severity > 0 (%) is calculated among distinct village-round observations in the corresponding regression sample, and Observations is the number of household-round observations. Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Appendix Table A5. Household welfare, village fixed effects

	(1) Continuous	(2) Discrete
Flood severity (# days)	-0.014***	
	(0.003)	
Low flooding		-0.081**
		(0.038)
Severe flooding		-0.238***
		(0.055)
Treatment × Flood severity	0.010***	
	(0.003)	
Treatment × Low flooding		0.016
		(0.050)
Treatment × Severe flooding		0.202***
		(0.070)
District × Round FE	Yes	Yes
Village FE	Yes	Yes
p: per-day loss fully offset	0.329	
p: low flood fully offset		0.061
p: severe flood fully offset		0.529
Observations	20375	20375

Notes: The dependent variable in both columns is the standardized overall household-welfare index, which summarizes the physical-health, coping, asset-damage, food-security, and agricultural-harvest domains reported in Table 6. Column (1) measures flood severity continuously as the number of synthetic-aperture-radar (SAR) flood days. Flood severity is the control-group association between an additional flood day and household welfare, and Treatment × Flood severity is the difference in that association between treatment and control households. Column (2) uses mutually exclusive no-, low-, and severe-flood categories. The low- and severe-flood coefficients are relative to no flooding among control households, and the corresponding treatment interactions are changes in the treatment-control difference relative to its value under no flooding. Treatment assignment is time-invariant within villages, so its main effect is absorbed by the village fixed effects and is not reported. Every regression absorbs district-by-round and village fixed effects. The three reported p-values test, respectively, whether the sum of the flood-severity coefficient and its treatment interaction equals zero, whether the sum of the low-flood coefficient and its treatment interaction equals zero, and whether the sum of the severe-flood coefficient and its treatment interaction equals zero. Standard errors are clustered by gram panchayat and reported in parentheses. Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Appendix Table A6. Ex ante adaptation components

Outcome	Treatment	Control Mean	N	Clusters
Advance-preparation index (all five domains)	0.054** (0.022)	-0.000	20,771	319
Human health-related	0.030 (0.024)	-0.000	15,188	319
Bought supplies (=1)	-0.001 (0.014)	0.288	4,912	319
Bought/arranged a boat (=1)	0.011 (0.009)	0.092	15,183	319
Health-related precaution (=1)	0.014 (0.010)	0.227	10,271	319
Stocked medicines (=1)	0.010 (0.009)	0.213	14,706	319
Stocked equipments to protect yourself from insects, etc (=1)	0.003 (0.009)	0.247	14,706	319
Stocked supplies for sanitary purposes (=1)	0.006 (0.008)	0.158	14,706	319
Animal-related	0.041* (0.023)	-0.000	20,760	319
Vaccinate livestock (=1)	0.006 (0.005)	0.030	5,577	319
Took precautions to keep animals safe (=1)	0.017 (0.011)	0.323	20,760	319
House / nearby environment	0.051*** (0.019)	0.000	20,765	319
Got a motor to drain water (=1)	0.007** (0.003)	0.036	20,760	319
Secured house (=1)	0.011 (0.007)	0.172	15,183	319
Secured land (not including agricultural land) (=1)	0.018*** (0.005)	0.070	15,183	319
Took measures to prevent water collecting (=1)	0.004 (0.005)	0.069	14,706	319
Cleaned around the house (=1)	0.006 (0.009)	0.271	14,706	319
Dug a drainage canal (=1)	0.010* (0.005)	0.081	14,706	319
Agricultural land / crop-related	0.031* (0.017)	0.000	20,766	319
Moved food/crops to higher ground (=1)	0.010 (0.011)	0.160	5,577	319
Took agriculture related precautionary measures (=1)	0.005 (0.004)	0.061	15,183	319
Decided to not hire labour (=1)	0.001 (0.001)	0.009	15,183	319
Land preparation adjustments	0.006 (0.003)	0.049	20,760	319
Seed-related adjustments	0.004 (0.004)	0.053	15,183	319
Fertilizer-related adjustments	0.004 (0.004)	0.050	15,183	319
Planting, sowing, nursery-related adjustments	0.005 (0.004)	0.053	15,183	319
Harvest-related adjustments	0.005* (0.003)	0.031	20,760	319
Irrigation-related adjustments	0.003 (0.002)	0.017	15,183	319
Food- and water-related	0.043** (0.020)	0.000	20,765	319
Stored drinking water (=1)	0.009 (0.006)	0.067	10,489	319
Stocked food (=1)	0.024** (0.012)	0.426	15,183	319
Ensured clean water for drinking or sanitation purposes (=1)	0.012* (0.007)	0.117	14,706	319
Purchased firewood/coal/cowdung/cakes/LPG etc. (=1)	0.004 (0.009)	0.233	14,706	319
Last-minute (rapid-response)	-0.008 (0.026)	-0.000	20,760	319
Moved household items (=1)	0.000 (0.009)	0.169	20,760	319
Household members moved (=1)	-0.005 (0.005)	0.059	20,760	319
Moved documents to a safe place (=1)	0.004 (0.008)	0.171	20,760	319
Moved animal fodder (=1)	-0.008 (0.008)	0.089	5,577	319
Moved food to higher ground (=1)	-0.009 (0.011)	0.184	10,271	319

Notes: The table decomposes the advance-preparation results in Table 4 into domain indices and the household survey items entering them. The 2025 questionnaire first asked, "Did your household undertake any measures to protect household members, household items, or livestock, or any other protective measures against flooding in 2025?" Households reporting measures were asked, for each reported month, "What steps did you take in [month]?" Follow-up prompts elicited the agricultural precautions ("What agriculture related precautionary measures did you take?") and health precautions ("What health related measures did you take?") listed in the indented rows. The first row is the advance-preparation index across five domains: human health, animals, house or nearby environment, agricultural land or crops, and food or water. Each domain index is followed by indicators for the displayed response options that enter it. The last-minute (rapid-response) index and its displayed items are reported separately and do not enter the five-domain advance-preparation index. Each row is a separate OLS regression of the listed outcome on treatment, absorbing district-by-round fixed effects and fixed effects for the observed number of SAR flood days. Standard errors are clustered by gram panchayat and reported in parentheses beside each coefficient. Control Mean is the unweighted mean of the outcome among control households in that row's estimating sample. Observations and Clusters are reported by row and can vary because the underlying questions were not asked of every household in every round. Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Appendix Table A7. Welfare components

Outcome	Treatment × No flood	Severe flooding	Treatment × Severe flooding	Control Mean	N	Clusters
Overall welfare index	0.001 (0.028)	-0.359*** (0.082)	0.299*** (0.100)	0.000	20,375	319
Physical health	0.019 (0.021)	0.008 (0.043)	0.053 (0.063)	-0.000	20,106	319
Illness subindex	0.002 (0.021)	-0.011 (0.053)	0.082 (0.071)	-0.000	20,106	319
Share of HH members sick (0-1)	0.003 (0.008)	-0.004 (0.021)	0.014 (0.026)	-0.396	20,103	319
Share of days any hh member was sick (0-1)	-0.000 (0.002)	0.000 (0.003)	0.005 (0.004)	-0.052	20,105	319
Illness treatment cost	-0.043 (0.695)	-0.618 (1.712)	2.522 (2.131)	-17.221	20,104	319
Injury subindex	0.008 (0.020)	0.009 (0.039)	0.033 (0.052)	-0.000	20,106	319
Share of HH members injured (0-1)	0.001 (0.004)	-0.002 (0.008)	0.008 (0.010)	-0.100	20,106	319
Share of days any hh member was injured (0-1)	0.001 (0.001)	0.001 (0.002)	0.002 (0.002)	-0.020	20,105	319
Injury treatment cost	-0.077 (0.210)	0.151 (0.566)	-0.065 (0.681)	-3.508	20,103	319
Any household member died	0.005 (0.003)	0.002 (0.008)	-0.003 (0.010)	-0.040	20,106	319
Coping	-0.008 (0.022)	-0.174** (0.078)	0.211** (0.089)	0.000	20,106	319
Consumption coping subindex	0.003 (0.023)	-0.018 (0.075)	0.106 (0.084)	-0.000	20,106	319
Postponed large purchases (=1)	0.003 (0.011)	-0.039 (0.031)	0.053 (0.035)	-0.331	20,106	319
Took out money from savings (=1)	0.002 (0.010)	0.031 (0.027)	0.000 (0.033)	-0.324	20,106	319
Reduced consumption of essential items (=1)	-0.002 (0.010)	-0.012 (0.029)	0.056* (0.033)	-0.364	20,106	319
Loan coping subindex	-0.018 (0.021)	-0.209*** (0.072)	0.245*** (0.077)	0.000	20,106	319
Loan Amount Since Aashadh (1K Rs.)	-2.129 (1.604)	-7.762 (4.837)	12.632** (5.064)	-36.345	20,099	319
Mortgaged assets (=1)	-0.002 (0.008)	-0.083*** (0.023)	0.076*** (0.026)	-0.164	20,106	319
Livestock coping subindex	0.001 (0.023)	-0.107 (0.066)	0.059 (0.073)	-0.000	20,079	319
Animals sold	-0.019 (0.016)	-0.044 (0.046)	0.029 (0.053)	-0.205	20,079	319
Animals rented	0.014 (0.010)	-0.039** (0.019)	0.016 (0.025)	-0.106	20,079	319
Asset damage	0.002 (0.030)	-0.370*** (0.106)	0.217* (0.119)	-0.000	20,375	319
Non-living damage subindex	-0.009 (0.034)	-0.464*** (0.113)	0.246 (0.151)	-0.000	20,375	319
Any document damaged	0.004 (0.005)	-0.021 (0.019)	0.021 (0.019)	-0.050	14,521	319
House damage cost	0.130 (0.268)	-2.347*** (0.526)	1.907*** (0.719)	-2.956	19,711	319
Items damage cost	-0.076 (0.129)	-1.035*** (0.339)	0.234 (0.538)	-1.104	19,853	319
Rebuild cost	-0.048 (0.227)	-2.744* (1.438)	1.970 (1.573)	-1.248	14,068	319
Home Flooded (=1)	-0.009 (0.015)	-0.237*** (0.048)	0.127** (0.059)	-0.147	20,363	319
Living-asset damage subindex	-0.022 (0.029)	-0.193** (0.095)	0.018 (0.105)	-0.000	20,079	319
Animals died	-0.026 (0.041)	-0.243* (0.128)	0.036 (0.141)	-0.536	20,079	319
Animals sick	-0.027 (0.040)	-0.264** (0.131)	0.026 (0.144)	-0.623	20,079	319
Human capital subindex	0.033 (0.024)	-0.041 (0.054)	0.146** (0.062)	-0.000	20,093	319
Share of HH members absent from school (0–1)	0.013 (0.010)	0.009 (0.025)	0.034 (0.029)	-0.303	11,704	319
Share of HH members absent from work (0–1)	0.004 (0.005)	-0.012 (0.013)	0.027* (0.014)	-0.182	20,091	319
Share of days absent from school, average member	0.001 (0.001)	0.001 (0.002)	0.002 (0.002)	-0.017	8,448	319
Share of days absent from work, average member	0.001 (0.002)	-0.004 (0.003)	0.007* (0.004)	-0.037	14,514	319
Food security	-0.004 (0.030)	-0.103* (0.061)	0.158** (0.074)	-0.000	14,524	319
You were worried you would not have enough food to eat (=	0.002 (0.013)	0.023 (0.026)	0.020 (0.033)	-0.285	14,524	319
You were unable to eat healthy and nutritious food (=1)	0.001 (0.013)	-0.063** (0.026)	0.086*** (0.031)	-0.394	14,524	319
You ate only a few kinds of foods (=1)	-0.012 (0.013)	-0.046* (0.023)	0.061** (0.031)	-0.420	14,524	319
You had to skip a meal (=1)	-0.001 (0.011)	-0.046** (0.022)	0.054* (0.028)	-0.197	14,524	319
You ate less than you thought you should (=1)	-0.009 (0.013)	-0.023 (0.024)	0.047* (0.028)	-0.309	14,524	319
Your household ran out of food (=1)	-0.001 (0.009)	-0.026 (0.021)	0.041 (0.026)	-0.171	14,524	319
You were hungry but did not eat (=1)	0.007 (0.010)	-0.054** (0.026)	0.067** (0.029)	-0.205	14,524	319
You went without eating for a whole day (=1)	0.000 (0.009)	-0.022 (0.018)	0.025 (0.023)	-0.132	14,524	319
Agricultural harvest	-0.008 (0.039)	-0.481*** (0.084)	0.283*** (0.100)	0.000	14,205	319
No crop loss from flooding (=1)	0.004 (0.016)	-0.196*** (0.046)	0.163*** (0.055)	0.643	14,205	319
Paddy harvest (kg, uncond.)	-12.314 (38.514)	-185.429*** (47.137)	31.737 (59.678)	366.983	14,151	319

Notes: The table decomposes the household welfare results in Table 5 into sub-indices and the individual household survey items that constitute them, using the discrete flood-severity specification. Each row is a separate regression of the listed outcome on indicators for low and severe flooding and on treatment interacted with each of the three flood-severity bins, absorbing district-by-round fixed effects, with standard errors clustered at the panchayat level and reported in parentheses beside each coefficient. Three of the five estimated coefficients are reported: the treatment effect among households experiencing no flooding, the level effect of severe flooding, and the treatment effect among households experiencing severe flooding. The low-flooding indicator and its interaction with treatment are estimated but not shown. Rows run from the overall welfare index through its sub-indices, with component items indented beneath the sub-index they enter. Control Mean is the mean of the outcome among control households in the estimating sample. Observations and clusters are reported per row and vary with item coverage. Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Appendix Table A8. Pooled OLS and IV household welfare index

	Household welfare index		
	(1) OLS	(2) 2SLS (IV)	(3) OLS
Treatment (ITT)	0.034		0.001
	(0.028)		(0.028)
Any alert = 1 (LATE)		0.373	
		(0.329)	
Any flooding (low or severe)			-0.226***
			(0.045)
Treatment × Any flooding (low or severe)			0.132**
			(0.055)
Kleibergen-Paap F		63.84	
District × Round FE	Yes	Yes	Yes
p: any flooding fully offset			0.027
Observations	20375	20338	20375

Notes: The dependent variable in all three columns is the standardized overall household-welfare index. Column (1) reports the intention-to-treat estimate from an OLS regression of the index on treatment assignment. Column (2) reports a two-stage least squares estimate in which an indicator for receiving any flood alert is instrumented with treatment assignment. The survey measured alert receipt by asking, "Did you receive any flood warning/alert?" Thus, the coefficient in column (2) is a local average treatment effect of alert receipt for households whose receipt responds to treatment assignment, rather than an intention-to-treat effect of assignment. The Kleibergen-Paap F statistic reports the strength of treatment assignment as an instrument for alert receipt. Column (3) is an OLS heterogeneity specification that pools low and severe flooding into Any flooding and uses no flooding as the omitted category. In that column the Treatment coefficient is the treatment-control difference among households experiencing no flooding, the Any flooding coefficient is the control-group difference associated with any flooding, and the Treatment × Any flooding coefficient is the differential treatment-control difference under any flooding relative to no flooding; the treatment-control difference under any flooding is therefore the sum of the Treatment and Treatment × Any flooding coefficients. The reported p-value tests whether the sum of the Treatment, Any flooding, and Treatment × Any flooding coefficients equals zero, i.e., whether treatment fully offsets the control-group welfare difference associated with any flooding. All three regressions absorb district-by-round fixed effects. Standard errors are clustered by gram panchayat and reported in parentheses. Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.