Direct Damage, Indirect Costs: The Incidence of Natural Disasters

Amen Jalal* Pol Simpson‡

July 14, 2026

Abstract

Natural disasters affect not only the locations they physically strike, but also surrounding areas through labour and product market linkages. As a result, the economic incidence of disasters differs from their physical footprint. We argue this in the context of Pakistan's 2022 floods, combining an original panel of 5,100 rural households with satellite and administrative damage data, and a re-centered hydrological instrument to separately identify the impacts of direct and indirect flood exposure. On its own, direct flood exposure causes physical damage but otherwise appears to leave households largely unharmed. Accounting for indirect exposure overturns this picture. A one-standard-deviation increase in surrounding flooding raises unemployment by 7 percentage points and local food prices by 9–18%, and lowers real consumption by 16%. All else equal, direct exposure made a household more likely to receive aid, while having flooded neighbours made it less likely. To ask how disaster aid should instead be allocated, we build a spatial general-equilibrium model, and find substantial welfare gains from targeting economic incidence rather than physical damage. Ignoring indirect exposure both understates the cost of large-scale climate shocks and misallocates the assistance that follows.

*London School of Economics (LSE); u.jalal@lse.ac.uk.

†London School of Economics (LSE); p.simpson1@lse.ac.uk.

‡We gratefully acknowledge funding from IGC, CID at Harvard University, CEPR STEG, and STICERD at LSE. We thank Clare Balboni, Oriana Bandiera, Gharad Bryan, Robin Burgess, Maarten De Ridder and Daniel Sturm for invaluable advice and support. For their outstanding research assistance, we thank Zahra Tajammal, Isra Imtiaz and Maheen Memon. For very helpful comments, we also thank seminar participants at the LSE Economics of Environment and Energy (EEE) hour and LSE Environment Camp.

1 Introduction

It is common to think of the costs of a natural disaster as borne by those it directly hits. Most disaster relief reflects this view, targeting easily verifiable physical damage to homes and infrastructure. Yet if there is trade, natural disasters (such as floods) can also affect surrounding areas through labour and product market linkages. For example, during a flood, households the water doesn't reach can suffer economic losses, whilst those it does may be partly insured if they have access to unaffected markets.

In this paper we show that indirect exposure, which we define as flooding in nearby areas, drives significant economic losses. Ignoring indirect exposure leads us to both underestimate the impacts of floods and overestimate the physical damage's share of the total loss. Our context is Pakistan's 2022 floods, which were caused by the heaviest monsoon in four decades and flooded 7% of the country. We assemble an original panel of 5,100 rural poor households, first observed before the floods and resurveyed one and two years later, and link it to satellite flood imagery and administrative damage records. To separately estimate the causal impact of direct exposure (flooding within 1km) and indirect exposure (1-15km), we use a recentered IV (Borusyak and Hull, 2023) based on a hydrological model. We find that direct exposure causes physical damage, whilst other economic costs, like lost employment and higher prices, are driven by indirect flooding. It follows that directly flooded households who are also surrounded by flooding are worst affected.

To ask how aid should be allocated in light of these results, we build a quantitative spatial model with trade and commuting. The model lets us derive the incidence of any hypothetical flood and in each case the optimal policy for a planner whose objective is to compensate loss. We then quantify the welfare gains from optimal targeting relative to targeting on physical exposure alone.

The core empirical challenge in the paper is estimating the causal effect of floods on household outcomes. Comparing households with high and low flood exposure may lead to biased estimates, since propensity to flood correlates with confounders like agricultural productivity. Our research design overcomes this challenge by comparing households with the same propensity to flood, who happened to experience different flood intensities in 2022 due to random variation in the exact pat-

tern of monsoon rainfall. We use a hydrological model to simulate expected flooding under many counterfactual rainfall realisations, subtract this from the hydrological models prediction for flooding under 2022 rainfall, and use this recentered prediction as an instrument for satellite-observed flooding in 2022. The approach is intuitively similar to controlling for a propensity score, except that we simulate it from our knowledge of how flooding is assigned, rather than estimating it from data. The identifying assumption is that the pattern of 2022 rainfall was as good as random with respect to unobserved drivers of household outcomes, after conditioning on district fixed effects and expected flooding. Separating direct from indirect exposure requires enough spatial variation in flooding to tell them apart, which we show holds in our setting.

With this approach in hand, we proceed to our empirical results. Restricting attention to the impact of direct flood exposure suggest the flood impacts were largely physical, and compensating damage to homes and infrastructure would be an adequate response. We find that direct exposure destroys homes and property, as expected. But on other economic outcomes – employment, food prices, savings, and credit – directly flooded households appear no worse off than the non-flooded, and on some measures they even appear to be better off.

That conclusion changes dramatically once indirect exposure is taken into account. While own flooding has no detectable effect on employment, households surrounded by one standard deviation (SD) more intense flooding become 7 percentage points (pp) more likely to be unemployed. This is driven by a decline in wage employment: the likelihood of being employed for wages falls by 17 pp, while the probability of self-employment rises by 9 pp. We interpret this shift as evidence of labour rationing, whereby some households displaced from wage work are pushed into lower-quality forms of self-employment, consistent with Breza et al. (2021).

One explanation for why directly exposed households are spared this fate lies in commuting patterns. We find that own flooding raises the share of men commuting more than a kilometre for work by 4 pp, as workers in flooded locations seek employment in less flooded markets elsewhere. By contrast, indirect exposure leaves commuting unchanged, since households surrounded by flooding have fewer nearby alternatives.

The same indirect channel prevails in product markets. A one-SD increase in surrounding flooding raises the price paid by households for meat by 9%, dairy by 18% and fruit by 17%. Direct exposure, in comparison, has no effect on prices once surrounding flooding is accounted for.

How do households cope with these losses? Those surrounded by flooding try to smooth the shock through conventional channels, but indirect exposure raises the probability of having tried and failed to obtain a loan by 4.9 pp, and new borrowing falls by roughly 15,600 rupees, consistent with credit rationing. Instead, they appear to turn to asset sales to make ends meet, though these estimates are imprecise. Direct exposure has no such effect on borrowing or asset sales. Consumption patterns tell a consistent story. Direct exposure alone appears to lower real expenditure by 5.9%, but this vanishes once surrounding flooding is accounted for. A one-SD increase in surrounding flooding is associated with a 16% decline in real expenditure.

Taken together, these results suggest that a flood's physical damage runs through direct exposure, while its economic costs – lost work, higher prices, tighter credit – operate through indirect exposure. A household hit on both margins, flooded and surrounded by flooding, is most affected.

Yet aid flows do not account for this pattern. Using household self-reports, we find that direct exposure raises the probability of receiving any flood assistance by 10.3 pp and food aid by roughly 4 pp. Indirect exposure, by contrast, *reduces* it – a one-SD increase in surrounding flooding lowers the likelihood of receiving food aid by 11.4 pp. This is consistent with either rationing within a local area, or aid providers being unable to reach locations surrounded by flooding.

To interpret these effects and ask how disaster assistance should be allocated across space in response, we build a spatial general-equilibrium model that combines Eaton-Kortum goods trade with commuting. Workers in the model live in one location, commute to a workplace in another, and buy goods produced anywhere. Thus, their welfare depends on their commuting and goods-market access. A flood does two things to the locations it hits: it lowers their productivity and destroys their housing. Because housing is a non-traded endowment, separable in utility, its loss falls entirely on the household whose home is damaged. Lower productivity, by contrast, raises the cost of goods and lowers wages, and in general equilibrium these changes propagate to prices and wages everywhere.

We calibrate the model's parameters from the literature and our own estimates, and simulate flooding in a stylised geography of identical locations. The model generates results consistent with our empirical findings: physical damage is caused entirely by direct exposure while changes in prices, labour-market outcomes, and welfare are driven disproportionately by indirect exposure.

Using the calibrated model, we ask how a fixed aid budget should be allocated across space. We consider a planner whose objective is to minimise the welfare loss remaining after aid, and who can choose a per-worker transfer for each location. We compare the optimal policy to two rules that target physical damage: a flat transfer to every flooded location, and an access-weighted transfer that falls with distance to dry land, reflecting the harder-to-reach pattern we find in the data.

The optimal policy delivers 8.7% more welfare per dollar than damage targeting, and 12.9% more than the access-weighted rule. It can improve on them in two ways: by varying transfers among the flooded, and by extending aid to unflooded locations. It gives most to the locations bearing the largest welfare losses, typically the interior of the flooded zone, far from dry land. So the share of loss compensated rises with the loss itself, where under a flat transfer it falls. The gains come from reallocating among the flooded at low budgets, and from reaching beyond the flood zone only at much larger ones.

How much targeting improves on damage-based aid depends on the disaster, and on three features in particular. The first is economic integration. When transport costs are high, each location is close to autarky, so a flood generates no spillovers and its incidence coincides with physical exposure — targeting on physical damages is then already optimal. Spillovers, and so the gains from targeting, arise only at the intermediate frictions of real economies. The second is the spatial concentration of the flood, which determines how many of the flooded can still reach dry land. In a dispersed flood most can, so the loss spills onto their dry neighbours and optimal aid reaches outside the flood zone. In a concentrated flood the interior cannot, so it bears the loss itself and optimal aid targets those trapped locations. The third is size. As a flood grows, the productivity loss spreads more evenly across flooded and unflooded locations alike, so their losses converge and physical exposure reveals less about who bears the loss. Together, this shows how much can be gained by targeting aid on economic incidence rather than physical exposure, and when those gains are largest.

By quantifying the relative importance of direct and indirect exposure, we contribute to a large literature on the microeconomic impacts of natural disasters, unifying two strands. The first estimates the impact of direct exposure on household outcomes, using households in physically unaffected locations as controls (Baseler and Hennig, 2023; Deryugina et al., 2018; Groen et al., 2020; Gallagher and Hartley, 2017; Gignoux and Menéndez, 2016; Kocornik-Mina et al., 2020; Boustan et al., 2020; Kirchberger, 2017; Patel, 2024). These designs implicitly estimate an average of direct and indirect impacts, in proportions that depend on the definition of the treatment unit and control group. We instead separate the two. A second strand shows that extreme weather events in one location have spillover effects on locations or firms not directly hit - through supply chain linkages (Barrot and Sauvagnat, 2016; Carvalho et al., 2021; Balboni et al., 2025), multinational production (Boehm et al., 2019; Castro-Vincenzi, 2024), evacuee labour supply (Fifer McIntosh, 2008) and food markets (Burgess and Donaldson, 2010; Allen and Atkin, 2022). We estimate direct and indirect effects jointly, and quantify the overall incidence of a disaster.

We also contribute to a small literature on the design of disaster assistance, by asking whether physical exposure should be a necessary or sufficient requirement for receiving assistance. There is very little normative work in this area - exceptions include Pople et al. (2026) on the optimal timing of disaster assistance, Fu and Gregory (2019) on reconstruction subsidies with rebuilding externalities, Tarquinio (2023) on rules versus discretion, and Gordon et al. (2025), who show that targeting aid on housing damage favours households who were wealthier to begin with, and who may therefore have better access to other means of smoothing.[1] We step back and ask who is affected in the first place, and provide an actionable method for targeting assistance on that basis. The model turns a flood map into a map of welfare incidence, which could be used to trigger payouts to individuals or localities, overcoming the challenges of observing indirect damage. More broadly, our work speaks to how private insurance, public insurance and other social protection can together help people navigate growing climate risks. Identifying who bears the loss lets us assess this system as a whole, and locate the gaps that new climate protection should fill.

[1] A related literature studies the targeting of transfers to the poor (e.g. Alatas et al., 2012, 2016; Hanna and Olken, 2018; Aiken et al., 2022; Banerjee et al., 2024). By contrast, we consider targeting based on loss after a specific event.

The rest of this paper proceeds as follows. Section 2 describes the 2022 floods and our study region. Section 3 introduces the data. Section 4 sets out the research design and Section 5 the first stage. Section 6 presents the results. Section 7 develops the spatial model, derives the optimal allocation of aid, and shows when targeting economic incidence improves most on physical exposure. Section 8 concludes.

2 Pakistan's 2022 Floods

Our setting is Pakistan's 2022 floods, which were the worst in the country's 75 year history.

Pakistan has a monsoon every summer, so heavy rain is not in itself unusual. In 2022, however, cumulative July–August rainfall across Pakistan was the largest in four decades: 114% above the 1981–2021 average, and 27% above the previous wettest year (2010) (Funk et al., 2015). The rain was concentrated in Sindh, the southernmost province and our study region, where it was more extreme still: 267% above the long-run average, and 44% above its own previous record (2003). It is very likely that climate change contributed to this unprecented event: a attribution study by World Weather Attribution estimated that the rainfall was 75% more intense than it would have been without human-induced global warming (World Weather Attribution, 2022).

The rain caused widespread flooding. Measured from MODIS satellite imagery (Nigro et al., 2014), 7.3% of normally dry land in Pakistan was inundated at some point during the season and 20.4% of Sindh; in the six districts where we work the figure reached 41%. Because Sindh is low-lying and flat, the water had nowhere to drain. Large areas stayed inundated for up to 8 months until the floodwaters evaporated.

The floods were extreme not only by Pakistan's own standards but globally. The Global Flood aabase of Tellman et al. (2021) maps, from daily satellite imagery, the extent of the 906 largest floods worldwide since 2000 – those logged in the Dartmouth Flood Observatory's global register of major floods and large enough for satellites to capture – recording for each the share of national land and population that went under water. Across these floods, the 90th-percentile event inundated 2.8% of national land and affected 1.7% of national population. Pakistan's 2022 floods covered 7.3% of

the country's land and affected 14% of its people, placing them at the 95th and 99th percentiles respectively.

This large-scale event is a compelling setting for our research question. In part this is a matter of power: the larger the event, the more likely we are to detect general-equilibrium impacts. But more importantly, enormous floods are worth understanding in their own right, since it is very likely that more of the floods we face in coming years will be large ones. Climate change is making weather extremes more frequent, and a single event increasingly covers more of a country at once.[2] We return to the policy implications of growing disaster scale in Section 7.

3 Data

To estimate the impact of direct and indirect flood exposure, we focus on a single event and exploit cross-sectional variation in flooding around each household. Our data come from three main sources: an original, geocoded panel of rural households; satellite and hydrological data that we use to measure flooding and construct our instruments; and administrative records of property damage. Table 3.1 summarises them.

3.1 Household survey data

We collect panel data on 5,100 poor, rural households in Sindh surveyed one and two years after the floods. The sample is drawn from a 2017 census of all low-income households in six districts of the province (Qambar Shahdadkot, Larkana, Matiari, Tando Muhammad Khan, Tando Allahyar and Sujawal). The data was originally collected by a local NGO to determine eligibility for development and poverty alleviation programmes and contains the address of the household, a list of household members (including their age, sex, education and occupation), household asset holdings across 21 asset classes and a poverty score constructed from these asset holdings. We use these data as a

[2]The area a single event covers has widened across hazards, most clearly for heat (Skinner et al., 2025) and drought (Brunner and Chartier-Rescan, 2024), and also for tropical-cyclone rainfall (Qin et al., 2024) and European floods (Kemter et al., 2020). In South Asia, both widespread monsoon rainfall (Roxy et al., 2017) and large-scale heat extremes (Banerjee et al., 2025) have become more common.

Table 3.1: Data sources

Data Type	Source	Years	Resolution
Panel A: Flood Exposure Measurement			
Flood extent (optical)	MODIS MCDWD	2022-23	250m, daily
Flood extent (radar)	Sentinel-1	Aug–Dec 2022	20m, multiple dates
Topography	FABDEM	Static	30m
Panel B: Instrumental Variables			
Rainfall	NASA IMERG	1998-2025	0.1°, daily
Topography	FABDEM	Static	30m
Panel C: Household and Damage Data			
Baseline census	NGO admin records	2017	GPS coordinates
Follow-up surveys	Primary collection	2023-24	GPS coordinates
Village census	Pakistan (Mouza) census	2020	Village level
Property damage	Urban Unit (Sindh govt)	2022-23	Structure level
Building footprints	VIDA (Google/Microsoft)	Static	Building level
Panel D: Agricultural output			
Vegetation greenness (NDVI)	MODIS MOD13	2001–2025	500m, monthly
Cropland mask	ESA WorldCover	2021	10m

sampling frame, as baseline controls to improve precision, and merge them with a 2020 village-level census recording local infrastructure (roads, markets, primary schools) for balance tests. Whilst restricting attention to the rural poor limits our ability to comment on the experience of wealthier or urban households, it's key benefit is that we can compare very similar households that happened to face very different flooding.

We designed the sample to span the range of flood exposure, selecting survey locations across the region and then randomly drawing households within them based on their pre-flood location. We surveyed households in two waves, running July–September 2023 and July–September 2024. With the assistance of local community members, we located 92% of households in 2023 or 2024; of those, almost all were still living in their 2017 village.[3] Attrition from the panel is low and balanced on flood exposure. In total we have 5,120 households and 33,202 people in 2023, and 4,902 households and 33,645 people in 2024. We conduct our main analysis on the panel of households surveyed in both waves.

[3]We had expected substantial permanent migration, but did not find this to be the case in our sample. Anecdotally, households expected reconstruction aid conditional on staying in place, and the worst-hit areas are not normally flood-prone, leaving little reason to relocate to reduce future risk.

The surveys cover damage to property, livestock and productive assets; lost income, savings drawdown, asset sales and new borrowing; labour supply, wages, occupational change, migration and commuting; and educational and nutritional investments — alongside self-reported flood exposure, including its duration, depth, and the household's evacuation decisions.

3.2 Administrative damage data

After the floods, the Government of Sindh conducted an in-person census of damage to residential property, recording 1.96 million damaged buildings (1.37 million of them fully destroyed) across the province. From this we construct small-area damage estimates on a standardised hexagonal grid, each cell of radius ~460m (resolution 8 of Uber's H3 system). For each hexagon we count the damaged buildings and divide by the total there, taken from the VIDA building-footprint layer, which fuses the Google and Microsoft open-buildings datasets. This yields, for each hexagon, a measure of physical damage independent of survey self-reports.

3.3 Satellite data to proxy agricultural output

We proxy agricultural output using satellite data measuring how green the land is. Specifically, we use the Normalised Difference Vegetation Index (NDVI), a widely used proxy for crop growth (Patel, 2024). We take monthly NDVI composites from MODIS at 500m, average them within each hexagon, and keep the hexagons that are predominantly cropland according to the ESA WorldCover land-cover map. A fall in NDVI after the flood then measures lost crop cover.

3.4 Flooding and hydrological data

Our flood and hydrological data come from four sources, summarised in Table 3.1. We measure flood extent and duration from two complementary satellite products: MODIS, an optical sensor that maps standing water at 250m but cannot see through cloud, and Sentinel-1, a radar sensor that sees through cloud and at night but revisits any given location less often. We measure depth by

combining the Sentinel-1 water maps with FLEXTH. Intuitively, at the edge of a flood the water surface sits at ground level, so the flood boundary reveals how high the water rose, and the depth at any interior point is that water level minus the ground elevation. To construct our instruments we use the LISFLOOD hydrological model, which takes daily rainfall and topography as inputs. Our rainfall data is satellite-based, from NASA's IMERG product, since Pakistan has few ground weather stations. For topography we use FABDEM, a bare-earth elevation model that strips buildings and forest canopy from the Copernicus terrain data, leaving the ground surface that routes water — drainage channels and irrigation canals. It misses large-scale flood defences, but these are largely absent in our setting.

We describe how these data are combined into our exposure measures and instruments below.

4 Research Design

Estimating the causal effect of flooding on economic outcomes faces identification challenges that could bias results in either direction. Flatter areas are more flood-prone but also more agriculturally productive – leading us to underestimate flood impacts. Conversely, high flood-risk areas may already be poorer due to past climate shocks or selective out-migration of wealthier households – leading us to overstate flood impacts.

We would avoid both biases if we could compare two locations equally prone to flooding, one of which purely by chance flooded in 2022 whilst the other did not. The insight from Borusyak and Hull (2023) is that we can construct exactly this comparison, because the selection driving flooding has a known source. How much a location floods is determined by the rainfall it receives and where it sits in the hydrological network. Knowing this, we can separate a location's underlying flood-proneness from the part of its flooding driven by the year's particular rainfall.

In practice, we use a hydrological model to simulate how much each location would have flooded under many alternative monsoon rainfall patterns. Averaging over these simulations gives each location's expected flooding, whilst the simulation using 2022 rainfall gives its predicted flooding for that year. We instrument satellite-measured flooding with the 2022 prediction, controlling for

expected flooding (which plays the role of a propensity score, computed exactly from the model rather than estimated from data). What remains is variation driven by the idiosyncratic rainfall of 2022: exactly the comparison we wanted.

4.1 Definition of flood exposure

We observe a sample of households at known locations, and for each we want to measure the flooding it faced – both at its own location and in the locations around it. To do this we lay a hexagonal grid over the study area and use satellite data to measure how badly each hexagon flooded. We then aggregate this into two measures of exposure for every household: *direct* exposure, the flooding in a small area around the household, and *indirect* exposure, a distance-weighted average of the flooding beyond that.

To measure flood intensity at each hexagon, we construct an index. It combines two satellite data sources, which offset each other's measurement limitations, and covers three dimensions along which floods differ – extent, duration, and depth. We build it using principal components analysis, which extracts the flood's underlying intensity as the common factor running through these measures.

The index draws on two satellite sources and one depth measure. MODIS images the surface every day but is often obscured by cloud, of which the monsoon brought plenty; Sentinel-1 sees through cloud and gives high-quality observations, but revisits any given place only every 6–12 days. Using each source we compute the mean share of land under water in each hexagon for each month from July to December. From the Sentinel-1 data and the FLEXTH algorithm described above we estimate maximum depth in each cell over the same period. This gives thirteen flood measures per hexagon: six monthly flood-shares from each satellite, and one depth. We combine the thirteen measures into their first principal component. Its weights are all positive and concentrated on the peak months, with depth contributing about as much as a single mid-event month.

This gives a flood-index value for every hexagon in the study area. Since we wish to test whether outcomes depend only on a household's own flooding, or also on flooding in the surrounding area

through its effect on local markets, we construct two measures of exposure. Write $Flood_j$ for the flood index at hexagon j and d_{ij} for the distance from household i to hexagon j.

Direct exposure D_i is the average index over the hexagons within a radius r of where the household lived, each weighted equally:

$$D_i = \frac{1}{n_i} \sum_{j\,:\,d_{ij} \le r} Flood_j, \qquad n_i = \#\{\, j : d_{ij} \le r \,\}. \tag{1}$$

Indirect exposure S_i is a weighted average of the index over hexagons beyond that neighbourhood, with weights that decay with distance:

$$S_i = \frac{\sum_{j\,:\,d_{ij} > r} e^{-(d_{ij} - r)/\lambda}\, Flood_j}{\sum_{j\,:\,d_{ij} > r} e^{-(d_{ij} - r)/\lambda}}. \tag{2}$$

The measure has two parameters, the own radius r and the decay scale λ (which we summarise by its half-mass h, the distance within which half the total indirect weight falls). Our results show sensitivity to both the direct exposure cutoff and the weights in the indirect exposure measure.

In our main specification we set the cutoff radius for own exposure at approximately 1.1km. For indirect exposure, we set the decay parameter at $\lambda \approx 4$ which implies that a hexagon at 4km has half the weight of a hexagon just beyond 1.1km, and that half of the weight of the indirect exposure measure is on hexagons within $\approx$ 7km of the household. Figure 9 in the appendix plots the resulting weights against distance.

4.2 Recentered IV approach

We would like to estimate the impact of direct and indirect exposure to flooding on household economic outcomes.

$$Y_{i,t} = \beta_D\, D_i(\mathbf{Flood}_{2022}) + \beta_S\, S_i(\mathbf{Flood}_{2022}) + \delta_d + \varepsilon_{i,t}, \tag{3}$$

where $\mathbf{Flood}_{2022}$ is the vector of flooding across all hexagons, $D_i(\cdot)$ and $S_i(\cdot)$ are the direct and indirect exposure defined above, δ_d is a district fixed effect, and $\varepsilon_{i,t}$ collects the unobserved determinants of Y.

The challenge is that flooding is not randomly assigned, so D_i and S_i may be correlated with unobserved determinants of Y. We address this by making the process that determines D_i and S_i explicit, so that we can control for selection. Both depend on flooding in all locations. In stylised form, flooding in each location at time t depends on the rainfall it has received, at its own location and upstream, and on the hydrological network F, which determines where water flows and accumulates:

$$\mathbf{Flood}_t = \mathbf{r}_t + F\mathbf{r}_{t-1} + F^2\mathbf{r}_{t-2} + \ldots + \boldsymbol{\nu}_t = \sum_{l=0}^{\infty} F^l \mathbf{r}_{t-l} + \boldsymbol{\nu}_t \tag{4}$$

Here $\mathbf{Flood}_t$ is an $N \times 1$ vector of flooding across locations and $\mathbf{r}_{t-l}$ is the vector of rainfall l periods earlier. F is the $N \times N$ flow matrix: F_{ij} is the share of water that moves from location i to location j each period, set by topography and by drainage infrastructure such as canals and embankments. Its powers F^l carry rain from upstream to downstream with a lag, so the sum accumulates water as it travels through the network.

We simulate flooding with LISFLOOD-FP, which produces daily predictions of flood depth and extent over approximately 150 days, yielding over 300 potential predictors of satellite flooding per grid cell. Rather than selecting a single summary measure ad hoc, we use a random forest to find the combination of hydrological outputs that best predicts our satellite-measured flood index – the optimal instrument in the sense of Chamberlain (1987), following Borusyak and Hull (2025). To avoid the instrument being fitted to idiosyncratic noise in each household's satellite measure (Chernozhukov et al., 2018), we cross-fit: we partition the grid cells into five spatially contiguous folds and, for each fold, train the random forest on the remaining four before predicting for the held-out fold.

We write the resulting predicted vector of flooding under rainfall $\mathbf{r}$ as $g(f(\mathbf{r}))$, where the hydrological model $f(\cdot)$ – the empirical counterpart of F above – takes rainfall and topography to predicted flood

depth and extent, and the random forest $g(\cdot)$ combines these into the best predictor of our flood index.

Our instruments are the predicted direct and indirect exposure, each recentered by subtracting its value under historical average rainfall $\bar{\mathbf{r}}$:

$$\tilde{z}_i^{\text{direct}} = D_i\big(g(f(\mathbf{r}_{2022}))\big) - D_i\big(g(f(\bar{\mathbf{r}}))\big), \tag{5}$$

$$\tilde{z}_i^{\text{indirect}} = S_i\big(g(f(\mathbf{r}_{2022}))\big) - S_i\big(g(f(\bar{\mathbf{r}}))\big). \tag{6}$$

The counterfactuals hold $D_i(\cdot)$, $S_i(\cdot)$, $f(\cdot)$, and $g(\cdot)$ fixed and vary only rainfall $\mathbf{r}$. The subtracted terms, $D_i(g(f(\bar{\mathbf{r}})))$ and $S_i(g(f(\bar{\mathbf{r}})))$, are each household's expected direct and indirect exposure – the exposure it would face under typical rainfall.

4.3 Identifying assumption

Households differ in how flood-prone they are, and these differences persist across years – a place reliably floods more if rain typically falls heavily on it or on the catchment draining into it. Our expected-exposure control absorbs this persistent component.

Identification comes from the deviation of 2022's rainfall from the typical pattern. Our IV is valid under the assumption that this deviation is unrelated to the other, unobserved shocks driving household outcomes – that the places where rain fell unusually hard in 2022 did not also differ, for unrelated reasons, in ways that affect their outcomes, after controlling for district fixed effects and expected exposure. Because our instrument varies only with rainfall, unobserved determinants of realised flooding – embankment breaches, local drainage failures – do not threaten identification unless they are themselves correlated with the 2022 rainfall deviation. Formally, collecting district fixed effects and expected exposure in ω,

$$\mathbb{E}[\varepsilon_{i,t} \mid \mathbf{r}_{2022}, \omega] = \mathbb{E}[\varepsilon_{i,t} \mid \omega]. \tag{7}$$

This assumption is testable, and we report a battery of checks – whether rainfall shows spatial variation our controls miss, whether it correlates across years after conditioning on expected exposure, and whether the 2022 anomaly was concentrated in areas with particular economic characteristics.

4.4 Standard errors

We report Conley (1999) standard errors, which allow the error term to be spatially correlated up to a chosen distance cutoff. Intuitively, we want to permit correlation at the scale at which treatment is assigned. Our treatment is generated by a single rainfall shock propagating through the hydrological network, so the relevant scale is the distance over which the instrument itself remains spatially correlated. We read this off the instrument's correlogram, correlating each hexagon's instrument value with those of its neighbours at increasing distances, and set the cutoff at the distance where that correlation falls to zero. The appropriate 'cutoff' will vary across regressions depending on the sample and treatment definition and is reported in all results tables.

5 First stage

Separating the impact of direct and indirect flood exposure on household outcomes relies on variation both across and within households in the two types of exposure. This places power constraints on how wide we can cast the net of 'indirect expxosure'. Too narrow, and direct and indirect exposure are highly correlated for the same household, because flooding is spatially correlatd. Too wide, and indirect exposure varies little across households. In our main specification we define direct exposure as flooding in a 1.1 km buffer around the household, and indirect exposure as a weighted average of flooding from 1–15 km, with a kernel decay parameter of $\lambda \approx 4$.

Figure 1 plots each household's predicted direct and indirect flood exposure for this specification with (right) and without (left) district fixed effects. The correlation between direct and indirect exposure is 0.58 without district fixed effects and 0.63 with them.

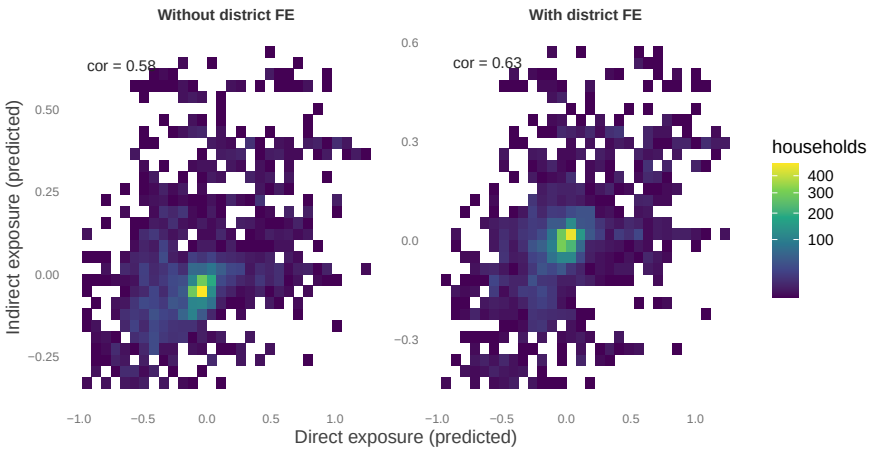


Figure 1: Joint density of predicted direct and indirect exposure

Notes: Bivariate histogram of household-level predicted (instrument) direct and indirect flood exposure, for the balanced panel of 4,683 households. Left: raw. Right: after residualising on district fixed effects.

Table 5.1 regresses the satellite exposure measure on both instruments – the recentered hydrological predictions. Each instrument strongly predicts its own measure: the own instrument loads predominantly on own exposure and the indirect instrument on indirect exposure, with partial F-statistics reported in the table.

Table 5.1: First stage

	Direct exposure (satellite)	Indirect exposure (satellite)
Direct exposure (Hydro IV)	1.31***	-0.03**
	(0.21)	(0.01)
Indirect exposure (Hydro IV)	-0.43*	1.30***
	(0.25)	(0.11)
Partial F on direct	37.6	6.5
Partial F on indirect	2.9	150.5
Conley cutoff (km)	14.3	14.3
N	4,683	4,683

Notes: First-stage regressions of each exposure measure (direct, indirect) on both instruments, the hydrological model's predicted flooding for the own and indirect areas. Partial F is the squared t-statistic on the matching instrument. Own-area radius $r \approx 1.1$ km, kernel half-mass $h \approx 7.1$ km. Balanced panel of households; district fixed effects; Conley standard errors. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

6 Empirical results

This section establishes three results. Direct, physical exposure to flooding causes physical damage, but it does not coincide with economic incidence - the losses to employment, prices, and consumption run through indirect exposure instead. Aid, however, tracks physical damage, and so misses much of where the loss falls.

6.1 Physical exposure to flooding causes physical damage

We begin by confirming what we would expect: flooding at a household's own location (direct exposure) destroys its property, whereas flooding nearby (indirect exposure) does not. For this we use the hexagon-level damage rate constructed from the administrative damage records, an objective measure of physical destruction within a small hexagon roughly 900 m across. Treating each hexagon as a small cluster of buildings, we ask how flooding in and around it affects the probability that a building is recorded by the government as flood-damaged. The sample is every hexagon in the union councils where we ran our surveys.

The first row of Table 6.1 regresses the damage rate on own exposure alone – flooding within 1.1 km of the hexagon. A one-standard- deviation increase in flood intensity within this radius raises the building-damage rate by 6.3 percentage points. The joint specification adds indirect exposure, controlling for flood intensity in the wider area. Its coefficient is small and far from significant ($t \approx 0.2$), and it barely changes the direct coefficient (6.6, against 6.3 on its own). Restricting to hexagons above the median building count, where the damage rate is measured more precisely, raises the own effect to 7.4 percentage points and leaves the indirect coefficient a precise null. Floodwaters more than a kilometre away do not destroy a house.

This may seem obvious, but that is exactly what makes it a useful test of our design. The chief concern it addresses is measurement error: if a household's own flooding is mismeasured, the flooding of its neighbours may proxy for the true local intensity, producing an apparent relationship between indirect exposure and outcomes where none exists (Huber, 2023). Physical damage is the one out-

Table 6.1: Damage placebo

	Any building damage	Any damage (above-median density)
Direct	6.29**	7.24***
	(2.54)	(2.39)
Direct	6.56**	7.38***
	(2.82)	(2.51)
Indirect	-0.65	-0.34
	(3.43)	(3.03)
Control mean (%)	38.9	36.4
N (hexes)	4,429	2,222
Conley cutoff (naive spec)	12.0	12.0
Conley cutoff (joint spec)	18.0	16.0

Notes. 2SLS of the res-8 hexagon building-damage rate (damaged buildings / VIDA buildings, capped at 1) on flood exposure; building-weighted, district fixed effects, Conley standard errors. Columns: all hexagons; hexagons above the median building count. The Direct coefficient is shown own-only (top) then jointly with Indirect (bottom). Coefficients are in percentage points, per one standard deviation of the respective exposure. Control mean: the lowest decile of predicted own flooding. $^{*}p < .1$, $^{**}p < .05$, $^{***}p < .01$.

come for which we can theoretically rule out the possibility of such an effect – it is hard to see how water several kilometres away could destroy a building, so an estimated indirect effect here would itself be cause for suspicion. Finding a precise null instead gives us confidence that where indirect exposure does move other outcomes, it reflects a genuine spillover effect rather than mismeasurement of own exposure.

Physical exposure also reduces agricultural output. Table 6.2 regresses satellite greenness on flood exposure at three points in the crop calendar, controlling for the hexagon's historical greenness in each month. Greenness measures realised crop cover, a proxy for agricultural output. Own flooding sharply reduces greenness in the flooded month, an effect that does not persist beyond the season; indirect exposure leaves greenness unchanged. The loss is concentrated in rice, while cotton falls by less but stays depressed into later seasons. Unlike housing damage, output may respond to flood conditioning in other locations (e.g. via farmer input choices) so we do not necessarily expect to see null effects of indirect exposure.

Table 6.2: Flood exposure and agricultural output (crop greenness)

	Sep 2022 (flood)	Feb 2023 (next Rabi)	Sep 2023 (next Kharif)
Panel A. All crops			
Direct (own only)	-0.136***	-0.013**	0.000
	(0.020)	(0.006)	(0.004)
Direct	-0.134***	-0.011	-0.001
	(0.021)	(0.007)	(0.004)
Indirect	-0.004	-0.007	0.003
	(0.012)	(0.013)	(0.010)
N	4,075	4,087	4,087
Panel B. Rice			
Direct (own only)	-0.166***	-0.007	-0.007*
	(0.014)	(0.007)	(0.004)
Direct	-0.164***	-0.002	-0.003
	(0.014)	(0.007)	(0.004)
Indirect	-0.005	-0.017*	-0.012*
	(0.006)	(0.010)	(0.007)
N	2,705	2,711	2,711
Panel C. Cotton			
Direct (own only)	-0.046***	-0.038***	-0.006
	(0.014)	(0.004)	(0.004)
Direct	-0.044***	-0.041***	-0.013***
	(0.011)	(0.006)	(0.003)
Indirect	-0.004	0.006	0.018***
	(0.008)	(0.009)	(0.001)
N	546	546	546

Notes. 2SLS of hexagon-level satellite greenness (MODIS NDVI, level) on flood exposure, for cropland hexagons in the survey union councils. Panel A pools all cropland; Panels B and C restrict to hexagons growing rice and cotton before the flood. Columns are three points in the crop calendar: the flooded Kharif (September 2022), the following Rabi (February 2023) and the next Kharif (September 2023). Each regression controls for the hexagon's 2016–2020 average greenness in that calendar month; the crop panels additionally control for its GAEZ irrigated attainable yield for that crop. District fixed effects, Conley standard errors. The Direct coefficient is shown own-only (top) then jointly with Indirect (bottom), each per one standard deviation of the respective exposure. $^{*}p < .1$, $^{**}p < .05$, $^{***}p < .01$.

6.2 Physical exposure does not coincide with economic incidence

Having established that direct exposure causes physical damage and indirect exposure does not, we now show that this pattern is very different for the other economic consequences of floods - their effect on labour and goods markets.

The floods arrived before the summer crop could be harvested and left fields under water for months, into the season when the next crop would normally have been sown. A shock of this timing and duration would likely have destroyed a substantial share of the year's harvest and depressed the demand for agricultural labour in flooded locations; in reducing output, it would also have lowered the supply of food from these areas.

Under autarky, these effects would be confined to the places the water reached. The fall in agricultural labour demand would lower wages and employment there; the loss of local output would raise local food prices. Physical exposure and economic incidence would coincide.

When villages trade, a productivity or supply shock in one place can be passed through to others: consumers can change where they buy their food, and workers can change where they sell their labour without moving home. Physical exposure and economic incidence therefore need not coincide. A household flooded at home but not in the surrounding area can lean on that area, drawing food and work from beyond the water's edge. A household flooded in the surrounding area but not at home can instead be on the receiving end, bearing costs passed through from the wider shock even as its own land stays dry. And a household flooded both at home and all around is closest to autarky: with no unaffected neighbour to share with, it bears the shock in full.

Labour Markets

We turn first to labour markets. Each outcome is measured at the household level, as an average across its working-age men.[4] Restricting attention to direct exposure alone (first row, Panel A, all households), we find no significant effect of own flooding on the probability that men are in

[4]We drop women from our sample because the female labour force participation in Pakistan is very low – around 20% (Jalal 2026; Bandiera, Jalal, Roussille 2026).

Table 6.3: Household male labour supply (shares of adult men)

	Any employment	Wage employment	Self employment	Commutes >1 km
Panel A. All households ($N = 4,478$)				
Direct	-0.00	-0.02	0.02	0.04**
	(0.02)	(0.03)	(0.02)	(0.02)
Direct	0.02*	0.03*	-0.01	0.05***
	(0.01)	(0.02)	(0.02)	(0.01)
Indirect	-0.07*	-0.17***	0.09***	-0.02
	(0.04)	(0.06)	(0.03)	(0.05)
Control mean	0.81	0.66	0.17	0.42
Panel B. Pre-flood agricultural households ($N = 1,830$)				
Direct	-0.02	-0.06*	0.04**	0.02
	(0.02)	(0.03)	(0.01)	(0.03)
Direct	0.03**	0.04**	-0.00	0.06***
	(0.01)	(0.02)	(0.03)	(0.02)
Indirect	-0.16***	-0.29***	0.12***	-0.11
	(0.05)	(0.06)	(0.04)	(0.08)
Control mean	0.83	0.50	0.34	0.27
Panel C. Pre-flood non-agricultural households ($N = 2,648$)				
Direct	0.01	-0.01	0.01**	0.06***
	(0.01)	(0.01)	(0.01)	(0.02)
Direct	0.01	-0.01	0.01**	0.03*
	(0.01)	(0.01)	(0.01)	(0.02)
Indirect	0.00	-0.00	-0.01	0.11
	(0.03)	(0.03)	(0.01)	(0.07)
Control mean	0.79	0.75	0.04	0.54
Conley cutoff (naive spec)	9.4	9.4	9.4	9.4
Conley cutoff (joint spec)	14.3	14.3	14.3	14.3
SD of Direct	1.00	1.00	1.00	1.00
SD of Indirect	0.72	0.72	0.72	0.72

Notes. 2SLS regressions of the share of adult men in each labour status (distance: mean over all adult men) on flood exposure, controlling for the number of adult men. Three panels: all households; pre-flood all-agricultural households; the remaining households with an adult man. Each panel's N is in its title; the control mean (lowest decile of predicted own flooding) is the last row. Both "Direct" lines and the "Indirect" line add controls chosen by post-double-selection lasso from the predetermined SUCCESS 2017 baseline; district fixed effects and the instruments are always kept. "Direct (own only)": own exposure alone. "Direct"/"Indirect": own and indirect jointly. Conley cutoffs and regressor SDs (shared across panels) are in the footer. Coefficients are per one standard deviation of own flooding at r=1.1 km. Conley standard errors. $^{*}p < .1$, $^{**}p < .05$, $^{***}p < .01$.

any employment, nor a significant change in the composition of their work. Direct exposure does, however, raise the probability of commuting more than a kilometre to work – beyond what we treat as the household's own location – by 4 percentage points per standard deviation of flood intensity, against a control mean of 42%.

Taken on their own, these estimates might suggest the floods left the labour market largely unaffected.

Turning to the joint specification, which additionally controls for indirect exposure, we find a very different picture. It remains the case that direct exposure has no significant effect on employment outcomes. What is now revealed, however, is that indirect exposure drives large and statistically significant effects: increasing average flood intensity in the surrounding area by the same one-standard-deviation increment reduces the probability of any employment by 7 percentage points, the net result of a 17-point reduction in wage employment – primarily daily labour in this setting – partly offset by a 9-point increase in self-employment.

The insignificant results in the first specification were masking effects operating through indirect exposure. Once these spillovers are taken into account, we find that flooding in the surrounding area has large effects on employment.

These findings, which will motivate our model later, strongly suggest that households whose village was under water are relatively well equipped to adjust to a localised labour-market shock in order to find work. By contrast, a household that was not itself flooded but lies close to an area of inundation can experience large and significant negative market effects. If your house is flooded, you do not necessarily lose your job; but if the wider area around you is flooded, you do. Households at the centre of large areas of inundation are hit twice – first as their house is destroyed, and then as they are unable to find work.

These effects show considerable heterogeneity by the type of work a household was doing before the floods. If we split households into those whose working-age men all worked in agriculture before the floods and those working in other sectors, we see that the results are driven almost entirely by the agricultural households (Panel B): for them, the same increment of surrounding flooding lowers

any employment by 16 percentage points and wage employment by 29, partly offset by a 12-point rise in self-employment. Among the rest, the effects on employment are essentially zero, their only response being a statistically insignificant 11-point rise in commuting beyond a kilometre (Panel C). This suggests not only that the floods were particularly devastating for agricultural jobs, but that there are large frictions preventing agricultural workers from moving to other occupations to smooth what appears to be a flood-induced reduction in labour demand – the spatial adjustment that non-agricultural workers evidently retain.

Goods Markets

Goods markets follow the same pattern as employment: physical exposure is once again a poor guide to economic incidence.

Our household surveys asked about consumption of four food groups over the past week – meat, dairy, fruit and legumes. For each, households reported both the quantity consumed, in kilograms, and its value, in rupees. Dividing value by quantity gives a geocoded consumer price for each good.

Table 6.4 shows that, among households reporting positive consumption of a given food, direct flood exposure has negligible effects on its price – the coefficients are small and insignificant – whereas indirect exposure raises the prices of meat, dairy and fruit. The same increment of surrounding flooding raises meat prices by around Rs 52 per kilogram (about 9% of the control mean), dairy by Rs 38 (18%) and fruit by Rs 34 (17%); the fruit effect is the most precisely estimated, significant at the 1% level, with dairy significant at 5% and meat at 10%, while legumes show no significant response.

In robustness checks, these effects strengthen as we place more weight on locations further from the household, suggesting they are driven by broader, more regional market conditions – though our power to detect such regional effects is more limited. This is consistent with lower trade costs, or greater market integration across locations, in goods than in labour.

To foreshadow our model, these results are consistent with a world in which each location produces very similar goods and is small relative to the regional economy: flooding in the local area alone is not enough to move prices, whereas more regional flooding can.

Table 6.4: Food prices (Rs/kg)

	Meat (Rs/kg)	Dairy (Rs/kg)	Fruit (Rs/kg)	Legumes (Rs/kg)
Direct	6.6	12.6*	13.5*	-4.7
	(9.7)	(7.3)	(7.0)	(3.5)
Direct	-8.5	1.1	3.3	-3.4
	(8.7)	(6.5)	(4.4)	(4.4)
Indirect	52.2*	38.3**	33.5***	-5.2
	(31.4)	(19.4)	(12.5)	(7.7)
Control mean	583.03	211.45	200.44	329.79
N	1,363	3,605	1,182	3,910
Conley cutoff (naive spec)	9.4	9.4	9.4	9.4
Conley cutoff (joint spec)	14.3	14.3	14.3	14.3
SD of Direct	1.00	1.00	1.00	1.00
SD of Indirect	0.72	0.72	0.72	0.72

Notes. 2SLS regressions of household food unit values (Rs/kg, winsorised at p99) on flood exposure, one column per good. Both panels add lasso-selected SUCCESS 2017 baselines; district fixed effects and the instruments are always kept. Top panel: own exposure only. Bottom panel: own and indirect jointly. Each good's *N* is the households buying it. Coefficients are per one standard deviation of own flooding at r=1.1 km. Conley standard errors. $^{*}p < .1$, $^{**}p < .05$, $^{***}p < .01$.

6.3 Aid tracks physical exposure, not indirect exposure

So far we have seen that physical damage is driven by direct physical exposure, whilst the effects on jobs and prices are driven by indirect exposure – falling both on directly exposed households that were themselves surrounded by widespread inundation, and on households that were not directly flooded but lived close to such areas. In other words, for outcomes other than property damage, the economic incidence of the flood does not coincide with where it physically hit.

Around the world, disaster relief very rarely formally targets indirect exposure to floods. Aid primarily follows physical damage, or else flows to areas the government designates as 'affected' – on the basis of opaque combinations of damage assessment and politics.

Using our survey data on self-reported receipt of disaster relief in the year after the flood, we show that this event was no exception. Directly exposed – physically flooded – households were 10.3 percentage points more likely to receive some form of formal flood relief (food, shelter, healthcare and so on), and 4.0 points more likely to receive food aid, against control means of 49 and 23 percent.

Table 6.5: Aid receipt

	Any flood relief	Any food aid
Direct	0.078***	0.006
	(0.024)	(0.020)
Direct	0.103***	0.040*
	(0.033)	(0.022)
Indirect	-0.084	-0.114**
	(0.058)	(0.056)
Control mean	0.49	0.23
N	4,318	4,683
Conley cutoff (naive spec)	9.4	9.4
Conley cutoff (joint spec)	14.3	14.3
SD of Direct	1.00	1.00
SD of Indirect	0.72	0.72

Notes. 2SLS regressions of household aid receipt on flood exposure: indicators for any flood relief and any food aid. Both panels add lasso-selected SUCCESS 2017 baselines; district fixed effects and the instruments are always kept. Top panel: own exposure only. Bottom panel: own and indirect jointly. Control mean: the outcome in the lowest decile of predicted own flooding. Coefficients are per one standard deviation of own flooding at r=1.1 km. Conley standard errors. $^{*}p < .1$, $^{**}p < .05$, $^{***}p < .01$.

Moving to indirect exposure, the picture is very different: not only does indirect exposure fail to attract aid, it appears to push in the opposite direction. More indirectly exposed households are less likely to receive food aid (an 11.4-point reduction, $p < 0.05$, against a base of 23%) and less likely to receive any flood relief (8.4 points, $p \approx 0.15$). These results are consistent with some form of rationing within a region – where, holding your own flooding fixed, more flooding around you means less assistance – or with a physical barrier to aid reaching you when the surrounding area is more flooded, reflecting anecdotal evidence that aid flows to 'easier to reach' locations. Among households that were themselves flooded, this means that as long as no one else around you is flooded you receive relief, but if you are flooded and also surrounded by flooding you may receive very little.

The pattern of flood relief after the 2022 floods not only failed to compensate for indirect exposure; it gave less to the households hit twice, by both direct and indirect exposure.

6.4 Direct exposure is not a sufficient statistic for flood-induced losses

What we have shown so far suggests that focusing on direct exposure alone is likely to give a misleading picture of welfare. The directly exposed suffered physical damage; to these households flowed relief, and, in subsequent years, the largest household-targeted recovery programme – a conditional cash transfer of around $1100 for home reconstruction. By contrast, the indirectly exposed - both those who were not themselves flooded and those who were flooded but surrounded by widespread devastation – suffered job losses and higher prices, for which they received no compensation.

Finally, we turn to how households fared, and to the ways they coped.

Consider direct exposure alone. Total expenditure shows no significant change, but the local price index rises, so real expenditure falls – by 5.9% (Table 6.8). On the balance sheet, the directly exposed appear better placed: 5 percentage points more likely to hold savings, around 8,500 rupees less indebted, and some 8,900 rupees stronger in net savings (Table 6.7). Taken at face value, these results suggest a shock that imposed only modest consumption losses on the directly exposed while leaving their financial position, if anything, stronger.

Separating direct from indirect exposure reveals that this picture conflates two distinct responses. None of the balance-sheet response is attributable to direct exposure, which on its own has no significant effect on savings, debt, or net position (although note that these asset measures exclude housing, and so do not capture the losses there). The effect on consumption is likewise weak: total expenditure is indistinguishable from zero, and real expenditure falls by an imprecisely estimated 1.2%.

The balance-sheet response is driven instead by indirect exposure, and it does not represent the improvement it appears to. It is the indirectly exposed who borrow less: indirect exposure lowers new borrowing by around 15,600 rupees (Table 6.7). This lower borrowing seems to reflect credit constraints rather than a lack of desire to borrow – indirect exposure raises the probability of having tried and failed by 4.9 percentage points, against a control mean of 47% (Table 6.6).

Table 6.6: Dissaving: asset sales, savings, and borrowing

	Sold HH assets	Sold productive assets	Spent savings	Took new loans	Tried to borrow, failed
Direct	0.017	-0.002	0.003	-0.035**	-0.003
	(0.013)	(0.019)	(0.021)	(0.017)	(0.025)
Direct	0.003	0.003	0.011	-0.004	-0.018
	(0.010)	(0.020)	(0.024)	(0.016)	(0.019)
Indirect	0.048	-0.014	-0.030	-0.104***	0.049*
	(0.030)	(0.073)	(0.058)	(0.029)	(0.029)
Control mean	0.11	0.24	0.68	0.61	0.47
N	4,683	4,683	4,683	4,683	4,683
Conley cutoff (naive spec)	9.4	9.4	9.4	9.4	9.4
Conley cutoff (joint spec)	14.3	14.3	14.3	14.3	14.3
SD of Direct	1.00	1.00	1.00	1.00	1.00
SD of Indirect	0.72	0.72	0.72	0.72	0.72

Notes. 2SLS regressions of dissaving indicators on flood exposure, all households. Sold productive / HH assets are coping strategies (8) and (10); sold either combines them; spent savings is (3); took new loans is (4); savings or new loans is (3) or (4); tried to borrow and failed is a separate question. Both "Direct" lines and "Indirect" add lasso-selected SUCCESS 2017 baselines; district fixed effects and the instruments are always kept. Top: own exposure only. Bottom: own and indirect jointly. Control mean: the outcome in the lowest decile of predicted own flooding. Coefficients are per one standard deviation of own flooding at r=1.1 km. Conley standard errors. $^{*}p < .1$, $^{**}p < .05$, $^{***}p < .01$.

Table 6.7: Savings, loans, and assets

	Has savings	Has loans	Asset index	Savings (000 Rs)	Loans (000 Rs)	Asset value (000 Rs)	Savings − loans (000 Rs)	Savings + assets − loans (000 Rs)
Direct	0.05	-0.03***	-0.05	0.54	-8.50**	-1.01	8.89**	7.42
	(0.03)	(0.01)	(0.06)	(0.63)	(3.52)	(5.46)	(3.70)	(5.49)
Direct	-0.00	-0.00	-0.01	-0.53	-3.61	3.02	3.21	6.34
	(0.02)	(0.02)	(0.05)	(0.48)	(4.03)	(10.86)	(4.07)	(11.19)
Indirect	0.18	-0.09*	-0.11	3.77*	-15.63*	-15.30	18.42*	1.16
	(0.12)	(0.05)	(0.13)	(2.11)	(9.17)	(28.96)	(10.04)	(36.54)
Control mean	0.20	0.77	0.17	3.65	95.31	124.83	-92.06	34.84
N	4,683	4,683	4,683	4,683	4,683	4,683	4,683	4,683
Conley cutoff (naive spec)	9.4	9.4	9.4	9.4	9.4	9.4	9.4	9.4
Conley cutoff (joint spec)	14.3	14.3	14.3	14.3	14.3	14.3	14.3	14.3
SD of Direct	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
SD of Indirect	0.72	0.72	0.72	0.72	0.72	0.72	0.72	0.72

Notes. 2SLS regressions of household balance-sheet outcomes on flood exposure, all households. Has savings / has loans are indicators and the asset index is the standardised ownership index; savings, loans (new borrowing), the valued movable assets (the ownership-index items priced at 2023 values, land excluded), and the net measures are winsorised levels in thousands of rupees. Both "Direct" lines and "Indirect" add lasso-selected SUCCESS 2017 baselines; district fixed effects and the instruments are always kept. Top: own exposure only. Bottom: own and indirect jointly. Control mean: the outcome in the lowest decile of predicted own flooding. Coefficients are per one standard deviation of own flooding at r=1.1 km. Conley standard errors. $^{*}p < .1$, $^{**}p < .05$, $^{***}p < .01$.

Unable to borrow, they draw on their assets instead, becoming 4.8 percentage points more likely to sell household durables – around 44% above a baseline of 11%, though this effect is not statistically significant. The directly exposed, by contrast, show no comparable coping. The net savings position of the indirectly exposed rises accordingly, by about 18,400 rupees, but only because the borrowing they sought was unavailable, not through any accumulation of wealth; they are more likely to hold savings largely because the proceeds of these sales are held as cash. Net worth – savings and movable assets net of loans – is unchanged. What appears as a strengthening of the balance sheet is more likely a household depleting its assets in the absence of credit.

Consumption losses, finally, are largest for the indirectly exposed. Indirect exposure lowers nominal expenditure by 7.4% and, together with higher local food prices, real expenditure by around 16% – substantially more than on the direct margin, with only the real-expenditure effect reaching significance, at the 10% level (Table 6.8). Much of the decline that own exposure alone attributed to direct flooding therefore ran through the surrounding area, just as on the balance sheet.

Table 6.8: Consumption

	Total exp	Food share of exp	Health share of exp	Price index	Real total exp
Direct	-0.023	0.014***	-0.006	0.038*	-0.059*
	(0.026)	(0.003)	(0.011)	(0.020)	(0.033)
Direct	-0.002	0.013***	-0.019*	0.009	-0.012
	(0.041)	(0.004)	(0.010)	(0.015)	(0.050)
Indirect	-0.074	0.001	0.049	0.097**	-0.159*
	(0.085)	(0.012)	(0.040)	(0.050)	(0.090)
Control mean	9.28	0.21	0.89	-0.40	9.67
N	4,683	4,681	4,681	4,536	4,536
Conley cutoff (naive spec)	9.4	9.4	9.4	9.4	9.4
Conley cutoff (joint spec)	14.3	14.3	14.3	14.3	14.3
SD of Direct	1.00	1.00	1.00	1.00	1.00
SD of Indirect	0.72	0.72	0.72	0.72	0.72

Notes. 2SLS regressions of household consumption on flood exposure, all households. Columns: log total expenditure; the food and health expenditure shares; the food price index; and price-deflated real total expenditure (log total expenditure minus the price index). Both "Direct" lines and "Indirect" add lasso-selected SUCCESS 2017 baselines; district fixed effects and the instruments are always kept. Top: own exposure only. Bottom: own and indirect jointly. Control mean: the outcome in the lowest decile of predicted own flooding. Coefficients are per one standard deviation of own flooding at r=1.1 km. Conley standard errors. $^{*}p < .1$, $^{**}p < .05$, $^{***}p < .01$.

7 A model of disaster incidence

The empirical results show that the incidence of a flood departs from its physical footprint, and that aid does not track it. How, then, should aid be allocated across space? To answer this we develop a spatial general-equilibrium model, and use it to evaluate counterfactual aid policies.

Setup. We model a closed economy with N locations, inhabited by workers, landowners and firms. Each location i has a fixed residential population of L_i workers, who supply labour in the location of their choice. Workers derive utility from consuming goods and housing services, the latter from a non-tradeable endowment of housing. Each location also has a single landowner who owns the land, and consumes goods bought with land rents.

Goods market. The goods market builds on Eaton-Kortum, with a few departures. There is a continuum of goods, which can be produced anywhere. Each good is produced from labour and land using Cobb-Douglas technology with labour share α. Regions differ in their average productivity and their idiosyncratic productivity for each good. Productivity for good u in location j is $z_j(u)$ – drawn from a Fréchet distribution with scale parameter T_j and shape parameter θ_T. T_j captures the absolute advantage of a location, whilst θ_T governs dispersion across goods. The unit cost of producing good u in location j is:

$$\frac{c_j}{z_j(u)} \quad \text{with} \quad c_j \propto w_j^{\alpha} r_j^{1-\alpha} \tag{8}$$

where w_j is the wage rate and r_j is the land rental rate. Firms operate in perfect competition, so this will also be the mill price of good u in location j.

Workers and landowners have identical CES preferences over the continuum of goods, with elasticity of substitution σ. They can buy goods from anywhere, but shipping is subject to iceberg trade costs $\kappa_{ij}^G = \exp(\tau^G d_{ij})$, so that κ_{ij}^G units must be shipped from j for one to arrive in i. They source each good from the least-cost location, and the resulting share of spending by i on goods from j is:

$$\pi_{ij}^G = \frac{T_j \left(c_j \, \kappa_{ij}^G\right)^{-\theta_T}}{\sum_k T_k \left(c_k \, \kappa_{ik}^G\right)^{-\theta_T}}, \tag{9}$$

All else equal, i will buy more from j if it has higher productivity, lower wages and land rents, or is closer. The price index, or cost of living, in i falls with its goods-market access: $P_i \propto \text{GMA}_i^{-1}$, where $\text{GMA}_i \equiv \left[\sum_k T_k (c_k \kappa_{ik}^G)^{-\theta_T}\right]^{1/\theta_T}$.

Labour market. Workers reside in i but choose where to work. All else equal, workers prefer locations that pay higher wages w_j and have lower commuting costs $\tau_{ij} = \exp(\tau^L d_{ij})$. They also draw an idiosyncratic taste for each workplace ζ_{ij} from a Fréchet distribution with shape parameter θ_L. Putting this together, the share of i's residents who commute to j is

$$\pi_{ij} = \frac{(w_j/\tau_{ij})^{\theta_L}}{\sum_k (w_k/\tau_{ik})^{\theta_L}}, \tag{10}$$

rising in the wage on offer and falling in distance, with θ_L governing how much workers sort on the financial benefits of jobs rather than idiosyncratic preferences. The expected wage is $\bar{w}_i = \sum_j \pi_{ij} w_j$. Their expected utility $\bar{v}_i$ is higher, reflecting the value of choosing their preferred job. Analogous to the price index, it rises with commuting market access $\text{CMA}_i \equiv \left[\sum_j (w_j/\tau_{ij})^{\theta_L}\right]^{1/\theta_L}$, so that $\bar{v}_i \propto \text{CMA}_i$.

Factor markets and incomes. Firm revenue in location j is total spending on its goods, $X_j = \sum_i \pi_{ij}^G E_i$, where E_i is income in i and π_{ij}^G its expenditure share on j. We assume factor markets are perfectly competitive, so revenue is split in fixed Cobb–Douglas shares: workers earn $w_j L_j^w = \alpha X_j$, where $L_j^w = \sum_i L_i \pi_{ij}$ is the workforce commuting in, and the landowner earns $r_j H_j = (1-\alpha) X_j$.

Welfare. A worker spends labour income on goods and enjoys htheir housing endowment. Their utility is Cobb–Douglas in consumption and housing services (since housing is a non-traded endowment, all of the worker's income is spent on consumption goods). The expected utility of a worker resident in i (averaging over their idiosyncratic workplace preference draws) is therefore

$$V_i^w = \left(\frac{\bar{v}_i}{P_i}\right)^{1-\beta} B_i^\beta \;\propto\; \left(\text{CMA}_i \cdot \text{GMA}_i\right)^{1-\beta} B_i^\beta, \tag{11}$$

where $\bar{v}_i$ is the option value of the jobs they can reach, P_i their cost of living, B_i their housing services and β the weight of housing in utility. Worker welfare rises with market access: to well-paying jobs

(CMA$_i$) and to cheap goods (GMA$_i$). The landowner has no housing and consumes only goods, so their expected utility is $V_i^r = (1-\alpha)X_i/P_i$.

Competitive equilibrium. Given location fundamentals $\{T_i, H_i, B_i, L_i\}_{i=1}^N$, parameters $\{\alpha, \beta, \theta_T, \theta_L, \sigma\}$, and bilateral trade and commuting costs $\{\kappa_{ij}^G, \tau_{ij}\}$, a competitive equilibrium is a vector of wages $\{w_j\}_{j=1}^N$ at which agents solve the problems above – pinning down c_j, π_{ij}^G, P_i, π_{ij}, $\bar{w}_i$ and $\bar{v}_i$ – and all markets clear:

(i) Labour: The wage bill equals labour's share of revenue, $w_j L_j^w = \alpha X_j$, where $L_j^w = \sum_i L_i \pi_{ij}$ is the workforce commuting in.

(ii) Land: Rents are pinned by land's share of revenue, $r_j = (1-\alpha)X_j/H_j$.

(iii) Goods: Revenue equals spending on each location's output, $X_j = \sum_i \pi_{ij}^G E_i$, where income is $E_i = L_i \bar{w}_i + (1-\alpha)X_i$.

We take the geometric mean of wages as the numeraire, normalising it to one, $\prod_j w_j^{1/N} = 1$.

Flooding in the model. We define a flood as a simultaneous negative shock to productivity and housing. Locations which are flooded experience a reduction in T_j, lowering average productivity, and B_j, as houses are destroyed.

In the model, destruction to housing has no general equilibrium consequences. Housing is non-traded across households or locations – so there is no way that households can act to change their access to housing services in response to damage – and it is separable in utility, so it does not affect their labour supply or consumption decisions. In practice, housing damage may have wider consequences, through reconstruction demand and building-supply prices, but we omit these from the model for now, since the large reconstruction housing programme was only just beginning in 2024, the last time we surveyed our households.

Changes in productivity, on the other hand, can have impacts that ripple out across the entire economy. To see this, recall that the price in a location is pinned down by unit cost. Let p_j be the mill price of an average good produced in j. Firms are perfectly competitive, so they price at marginal

cost – the input bundle c_j per unit of output:

$$p_j \;\propto\; \frac{c_j}{z_j}, \qquad c_j = w_j^{\alpha} r_j^{1-\alpha}, \tag{12}$$

A flood lowers T_j, hence z_j: it takes more land and labour to make a unit of output, so unit cost rises. Differentiating, and splitting the bundle cost into its two factors,

$$\underbrace{-\,a_j}_{\text{unit cost up}} \;=\; \underbrace{d\ln p_j}_{\text{buyers pay more}} \;-\; \underbrace{\alpha\, d\ln w_j}_{\text{workers earn less}} \;-\; \underbrace{(1-\alpha)\, d\ln r_j}_{\text{landowners earn less}}, \tag{13}$$

where a_j is the proportional change in average productivity, negative under the flood. Zero profits imply that a fall in productivity must be absorbed by some mix of a higher price paid by buyers, lower wages, and lower land rents.

Intuitively, the trade and commuting elasticities govern how the loss is divided between prices, wages, and rents. If the trade elasticity is low, the flooded location passes its lower productivity into a higher price, leaving wages and rents largely unchanged. If it is high, the price cannot move, and the split between wages and rents depends on the commuting elasticity. If labour is immobile, the two fall together. If labour is mobile, the wage is held up by workers leaving to find work elsewhere, and the residual loads onto land rent.

These changes occur at every production location at once. To turn them into the welfare of the people who *live* in a given location i, recall that residents shop and work in places other than where they live. As a result, a resident of i is exposed to prices and wages wherever they buy goods and supply labour. Linearising around the pre-flood equilibrium we can collect these exposures into a single expression.

$$d\ln V_i^w = \underbrace{\beta\, d\ln B_i}_{\text{housing}} + (1-\beta)\Big[\underbrace{\textstyle\sum_j \pi_{ij}\, d\ln w_j}_{\text{wage}} - \underbrace{\textstyle\sum_j \pi_{ij}^G\, d\ln p_j}_{\text{cost of living}}\Big], \tag{14}$$

The change in average welfare for a worker living at i depends on the change in their housing endowment $\ln B_i$, the change in wages in all locations $d\ln w_j$ weighted by the probability that they

commuted to each location prior to the flood $\pi_i j$, and the change in prices in all locations $d \ln p_j$ weighted by the share of their expenditure that was spent on goods from j prior to the flood. The wage and price changes are themselves general-equilibrium responses to the flood.

The landowner holds no housing and lives off rent, so two channels collapse to one:

$$d \ln V_i^r = \underbrace{d \ln r_i}_{\text{rent}} - \underbrace{\sum_j \pi_{ij}^G \, d \ln p_j}_{\text{cost of living}}, \qquad d \ln r_i = d \ln X_i, \tag{15}$$

Together these equations map any flood to its welfare cost at every location. We now calibrate the model to quantify them.

7.1 Calibration

We are in the process of quantifying the model using data from Pakistan. The numerical results that follow come from a simulation of a stylised economy. The economy consists of $N = 400$ locations arranged on a 20×20 grid. Locations are equally spaced on a torus (a donut shape), so the grid has no edges and every location faces the same distribution of distances to all other locations. We also impose that all locations are identical in endowments ($L_i = H_i = B_i = 1$) and baseline productivity ($T_j = T$). Together these imply that each location has identical market access in the baseline.

To parameterise the model we estimate some parameters from our data, take conventional values from the literature where available, and otherwise choose illustrative values for the simulation. The labour share in production is $\alpha = 0.55$, which we take from Sotelo (2020). Preferences are CES across goods with elasticity of substitution $\sigma = 1.1$, which we estimate from the HIES, and Cobb–Douglas between goods and housing with a housing share $\beta = 0.09$, which we estimate as the rural-Sindh rent share in the 2018–19 HIES. The two Fréchet shape parameters are the trade elasticity $\theta_T = 4$, which we take from Simonovska and Waugh (2014), and the commuting elasticity which we set to $\theta_L = 1.5$. Trade and commuting costs are iceberg and exponential in distance, with a common gradient which we set to $\tau^G = \tau^L = 0.30$. A flooded cell loses a fraction $a = 0.50$ of its productivity

and $b = 0.50$ of its housing, which are illustrative but will later be estimated from housing-damage and agricultural-output data.

7.2 Disaster aid in the model

Our objective is to compare real-world examples of disaster aid policy to the model-implied 'optimal' allocation. To do this we need two things. First, a model of who is harmed by the disaster, which we have outlined above. Second, we need to write down what we think the planner is trying to achieve. We can compare any allocation of aid to the distribution of welfare impacts from the model, but to quantify how much worse one allocation is than another we need a way of aggregating individual outcomes into social welfare.

We assume the planner chooses a per-worker transfer y_i in each location to minimise the welfare loss that remains once aid is paid, subject to a fixed budget B:

$$\min_{\{y_i\}} \sum_i \omega_i L_i \left[u\left(V_i^{w,\text{pre-flood}}\right) - u\left(V_i^{w,\text{post-flood}} + y_i\right) \right] \quad \text{s.t.} \quad \sum_i L_i \, y_i = B, \qquad u(V) = \frac{V^{1-\varepsilon}}{1-\varepsilon} \tag{16}$$

where ω_i is the welfare weight on location i. Individual utility is linear in income, so the household's welfare V_i^w enters linearly. The planner, however, applies a concave transformation $u(\cdot)$, so that social welfare places more weight on households left worse off. Without this curvature the planner would be indifferent across allocations. Intuitively, the planner values each worker's flood loss in utility terms and allocates aid to minimise what remains of it. Since pre-flood welfare is sunk, this is equivalent to maximising post-aid social welfare.

This simple problem embeds several substantive assumptions.

First, the objective is to compensate households for what they lost, not to alleviate poverty. In principle, we allow pre-flood standing to enter the welfare weight ω_i, so the planner may value a given loss more when it falls on a poorer household. However, as a simplification we set $\omega_i = 1$, since our population of interest is rural and poor.

Second, the planner solves a static problem, allocating a fixed budget once the flood has happened. We ignore the dynamic impacts of assistance, which would require the planner to weigh the benefits of aid today against the cost of crowding out future private risk reduction (Baily, 1978; Chetty, 2006). Moral hazard would change the allocation derived here only if it were heterogeneous and positively correlated with damage, so that the worst-hit also changed their adaptation behaviour most in response to aid. Otherwise it changes only the total budget. In our context aid may even raise private risk reduction, both through income effects and by relaxing the credit constraints that stop poor households financing it.

Third, we ignore how the funds are raised. This is reasonable in our context, where most relief is externally financed by international donors and the affected population pays little tax.

Last, we assume aid takes the form of a lump-sum cash transfer with no general-equilibrium effects, so each household's welfare depends only on its own transfer, $V_i^{w,\text{post-aid}} \equiv V_i^{w,\text{post-flood}} + y_i$. The advantage is simplicity. With the loss distribution fixed by the flood, we can ask the core question directly – how well different targeting rules approximate the true distribution of losses – without the losses themselves shifting as we change the allocation. Were aid to have general-equilibrium effects, a transfer to one location would move the prices and wages faced by others, so every household's welfare would depend on the entire vector of transfers, $V^{w,\text{post-aid}} = V_i^{w,\text{post-flood}}(\mathbf{y})$. We leave the general-equilibrium effects of aid to future work.

7.3 Solving for Optimal Aid Policy

Given the planner's problem above, we solve for the optimal allocation. We consider two variants – one where every worker who receives aid gets the same size transfer $\bar{y}$ (the discrete optimal), and another where the transfer is allowed to vary across locations y_i (the continuous optimal). In both cases the transfer is the same for all workers who live in a given location.

The discrete case corresponds to policies such as the BISP cash transfer after the 2022 floods, a fixed per-household payment for which the government chose which districts were eligible. Here, the planner is allowed to choose the level of the transfer per worker $\bar{y}$ and which locations receive

it. Intuitively, we can think about the planner deriving the optimal policy by solving this problem in two steps. For each possible transfer size, the planner ranks locations by their marginal welfare gain, $u(V_i^{w,\text{post-flood}} + \bar{y}) - u(V_i^{w,\text{post-flood}})$, and, starting from those with the largest marginal gain, adds them to the list of eligible locations until the budget is exhausted. The optimal transfer size is then chosen by comparing the social welfare achieved by the optimal allocation at each transfer level. With common pre-flood welfare, as in our simulation, ranking by the marginal welfare gain is equivalent to ranking by the size of the loss; in general it is not, since the planner targets the worst-off rather than those who fell furthest.

In the continuous case the planner can give any non-negative amount to each location. The optimal policy will take the form of choosing a welfare cutoff $\bar{V}$, giving nothing to locations already at or above it, and topping the rest up to $\bar{V}$, with the cutoff set to exhaust the budget. Unlike the discrete case, this requires knowing not just who is worst off but by how much.

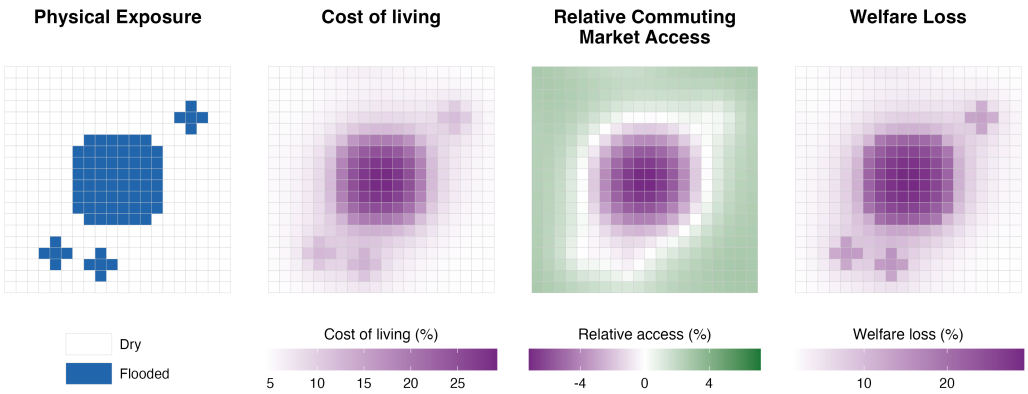
We contrast these two optimal policies against two stylised versions of the rules used in practice, both of which target physical damage rather than welfare loss. The first, damage targeting, pays the same amount to every worker in a flooded location. The second, access-weighted damage targeting, also pays only flooded locations but scales the amount down with distance to the nearest dry land, reflecting that remote locations are harder for relief to reach.

7.4 Welfare gains from optimal policy

We begin with a single flood, modelled as a discrete event – a location is either flooded or not. In this flood, 19% of locations (75 out of 400) were flooded. In each flooded location 50% of the housing endowment was destroyed and productivity fell by 50%. Since locations are symmetric in this economy, this implies that 9.5% of the whole housing stock was destroyed, and, in partial equilibrium, output in the economy would fall by 9.5%.

Figure 2 shows the flood (far left), the resulting welfare loss (far right), and the two general-equilibrium channels through which the flood reaches welfare: the cost of living residents face and their commuting market access. On average, across the whole economy, worker welfare falls by

Figure 2: Physical Exposure, Price Changes and Welfare Loss



Notes: From left: flooded locations (blue); the change in the cost of living residents of i face, $\Delta \ln P_i$ (the price index over the goods they buy); the change in their commuting market access, $\Delta \ln \text{CMA}_i$ (the option value of the workplaces they can reach); and worker welfare loss, $(V_i^{w,\text{pre-flood}} - V_i^{w,\text{post-flood}})/V_i^{w,\text{pre-flood}} \times 100$. Purple marks the adverse direction throughout. Average wage changes are not identified (since they are the numeraire) so panel 3 shows relative impacts on commuting access.

7.8% as a result of this flood. However, across locations there is a wide range in the welfare impact of this event - ranging from 29.2% for the most affected location to 1.4% in the least affected.

How does the welfare loss relate to physical exposure? For the average physically exposed worker this figure is 21.9% and for the average worker who was not flooded this figure is 4.5%. However, physical exposure is not a sufficient statistic for welfare loss - amongst the flooded the loss varies from 11.8% to 29.2%, whilst amongst the unflooded it varies from 1.4% to 12.9%, so that the worst-hit unflooded locations lose more than the least-hit flooded ones.

We showed above that flood losses depend not only on flooding in your own location, but on flooding everywhere, weighted by each location's importance as a trading or employment destination. We make this concrete by running the reduced form of Section 4.2 on the model's own output, regressing each location's outcomes on its direct and indirect exposure, built with the same kernel as in the data (Table 7.1). Physical damage to housing loads entirely on direct exposure: own flooding predicts it perfectly, and indirect exposure adds nothing. The economic outcomes are the reverse. On its own, direct exposure explains 80% of the variation in welfare loss and indirect exposure 94%; with both included, a standard deviation of indirect exposure raises welfare loss by 5.2 percentage points, against 2.9 for direct exposure. The same pattern runs through the two channels: a standard deviation

of indirect exposure raises the cost of living by 4.7 points and lowers relative commuting market access by 2.4, both far larger than direct exposure. The flood reaches welfare mainly through the prices households pay and the jobs they can reach, not the water at their own door.

Table 7.1: Direct and indirect exposure regression using model data

	Housing damage (%)	Cost of living (%)	Relative commuting access (%)	Welfare loss (%)
Panel A: direct exposure only				
Direct exposure	19.54	5.37	-2.07	6.81
R^2	1.00	0.74	0.64	0.80
Panel B: indirect exposure only				
Indirect exposure	14.81	6.08	-2.58	7.37
R^2	0.57	0.95	1.00	0.94
Panel C: both				
Direct exposure	19.54	1.80	-0.25	2.87
Indirect exposure	0.00	4.72	-2.39	5.19
R^2	1.00	0.99	1.00	1.00
SD (direct exposure)	0.39	0.39	0.39	0.39
SD (indirect exposure)	0.16	0.16	0.16	0.16
Control mean (least-exposed)	0.00	4.48	2.88	1.40
Locations	400	400	400	400

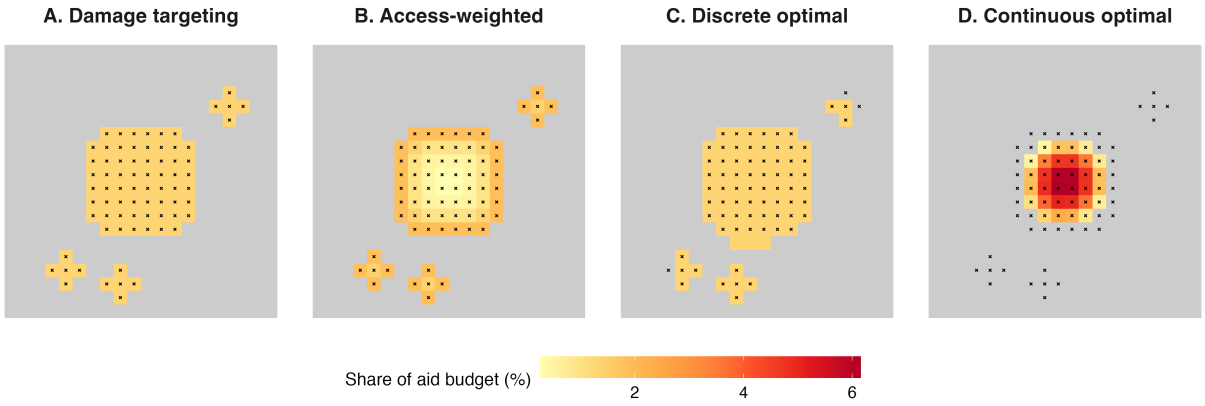
Notes: One observation per grid location. Each column regresses one outcome on standardised direct and indirect exposure, using the empirical kernel of Section 4.2; each coefficient is the percentage-point change in the outcome per one standard deviation of exposure. Panels A and B include each exposure alone; Panel C includes both. Housing damage is the share of the local housing stock destroyed; cost of living is $\Delta \ln P_i$; relative commuting access is $\Delta \ln \mathrm{CMA}_i$, measured against the economy-wide average (the numeraire fixes the average wage); welfare loss is in % of baseline. SD rows give each exposure's spread; the control mean is each outcome at the least-exposed location.

We now ask how a policymaker would allocate aid after this flood. We take the two optimal policies derived above, discrete and continuous, and compare them to the two status-quo policies, damage targeting and access-weighted damage targeting. We set the aid budget at 0.2% of pre-flood aggregate worker welfare, $B = 0.002 \cdot \sum_i V_i^w$, which equals 2% of pre-flood output at this calibration and keeps the budget consistent with the counterfactuals below; the welfare curvature is $\varepsilon = 1.5$.

Figure 3 shows the resulting allocation under each policy. Panel A is damage targeting; it mirrors the flood map exactly, giving every flooded worker the same transfer of 1.07% of baseline income. Panel B, the more realistic access-weighted policy, assigns transfers ranging from 0.20% to 1.50% of baseline income per worker: conditional on being flooded, a standard deviation more indirect

exposure – a location deeper inside the flooded zone, further from dry land – lowers the transfer by 0.37 percentage points of income. Panel C is the discrete optimal, which is nearly identical to damage targeting: at the same fixed transfer of 1.07% of income it covers 75 locations, swapping only the three least-affected flooded locations for the three worst-hit unflooded ones. Panel D is the continuous optimal, where the planner can vary the transfer across space. It is almost the mirror image of the access-weighted policy: transfers range from 0.42% to 4.93% of pre-flood income per worker, with the largest, 4.93%, going to locations at the centre of the main flooded zone, and nothing at all to 43 of the 75 flooded locations, on account of their proximity to unflooded land.

Figure 3: Aid allocation under each policy

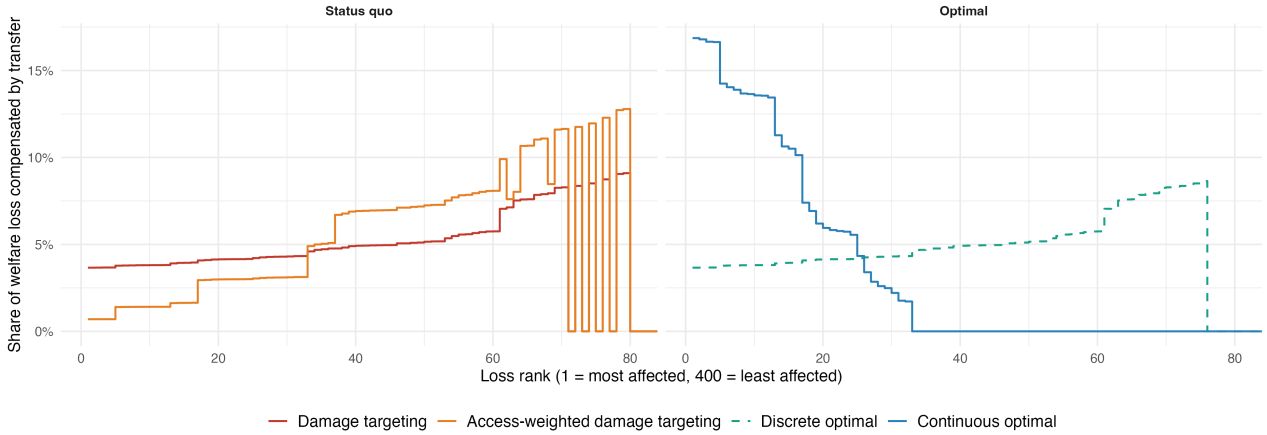


Notes: Each location's transfer as a share of the total aid budget, across the 400 grid locations; crosses mark flooded locations, grey cells receive nothing. A gives every flooded location the same transfer; B scales it down with distance to dry land; C is the discrete optimum; D is the continuous optimum.

These policies imply large cross-location variation in the share of loss that is compensated (Figure 4). A flat transfer, like the discrete optimal or damage targeting, gives the same amount to all, so it compensates a smaller share of loss for those hit worst – the worst-affected recover $\sim$ 4% of their loss, against $\sim$ 9% for the least-hit flooded locations. This is exacerbated under access-weighted damage targeting, since the worst-hit locations are the hardest to reach and so receive the least, recovering only $\sim$ 1%. The continuous optimal does the opposite, directing the largest compensation rate to the worst-affected – who recover $\sim$ 17% – focusing on insuring the most extreme losses.

We also compare the cost-effectiveness of assistance across these policies, in terms of welfare gained per dollar spent. Relative to damage targeting, the continuous optimal delivers 8.7% more welfare per dollar, and relative to access-weighted damage targeting, 12.9% more. The discrete optimal essen-

Figure 4: Share of welfare loss compensated by transfer, by loss rank



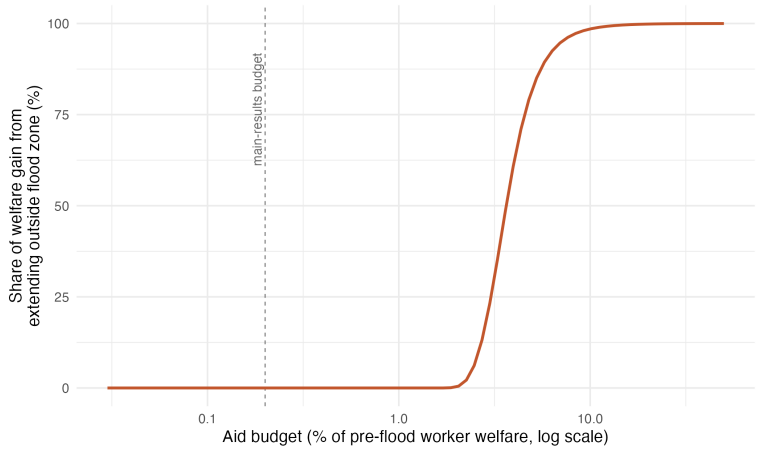
Notes: For each affected location ($\ell_i = V_i^{w,\text{pre-flood}} - V_i^{w,\text{post-flood}} > 0.01$), the share of loss covered by the transfer, y_i/ℓ_i, against the rank of ℓ_i from largest (rank 1) to smallest. Left: the two status-quo rules. Right: the two model-based rules.

tially ties damage targeting – the three-for-three swap adds almost nothing. The continuous rule's advantage comes from reweighting among flooded locations rather than from reaching unflooded ones, directing more to the deep interior and less to the periphery and outlying clusters that retain options on the dry side. Access-weighted damage targeting does the reverse, favouring the accessible periphery, which is why it does worst.

That the continuous optimal improves on the status quo by reallocating among the flooded, rather than by reaching outside the zone, is a feature of this budget. As the budget rises, the worst flooded losses are covered and the largest remaining losses fall outside the zone, so transfers to unflooded locations supply a growing share of the gain.

At a budget of 0.2% of pre-flood worker welfare, equivalently 2% of GDP and 2.6% of total welfare losses, the entire gain over the flat transfer comes from reallocating among the flooded (Figure 5). Sending resources outside the flood zone only begins to contribute once the budget covers about a third of total losses, and supplies half the gain once it covers roughly half.

Figure 5: Where the gains from optimal aid come from, by budget size



Notes: Share of the continuous-optimal policy's welfare gain over the flat status-quo transfer that comes from extending transfers outside the flood zone, rather than reallocating among the flooded, for the main-results flood, as the aid budget varies. The budget is on the horizontal axis as a percentage of pre-flood worker welfare (log scale); the dashed line marks the main-results budget. The remainder of the gain comes from reallocating among the flooded.

7.5 Determinants of the gains from targeting

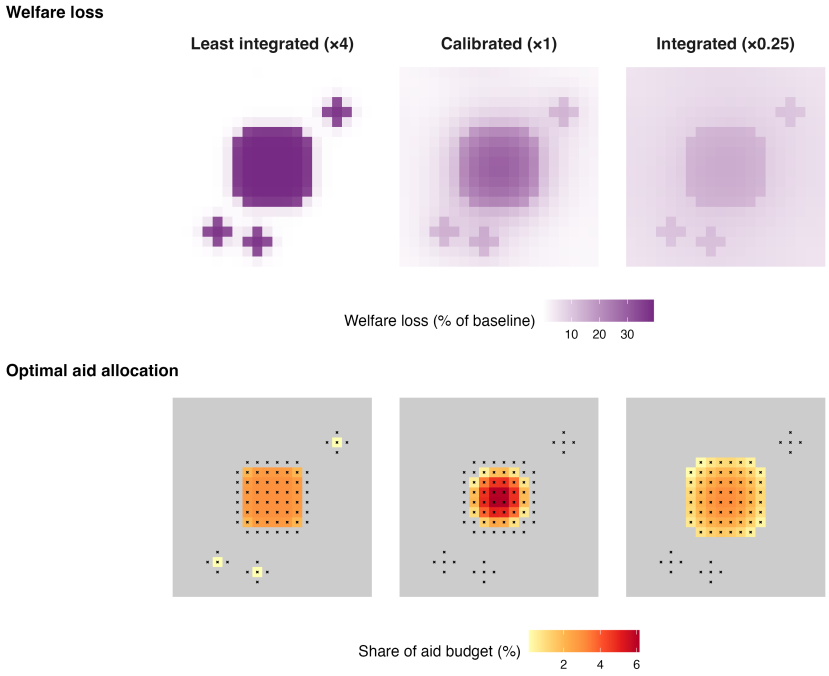
This illustrative example shows there is room to improve on status-quo aid policy when the incidence of the flood differs from physical exposure. The gains from targeting therefore depend on the gap between the two, which is shaped by three factors – trade costs, the spatial dispersion of the flood, and its size. We vary each in turn.

Transport costs

To vary economic integration, we scale the commuting and trade frictions together, from near-frictionless to near-autarky, and re-solve the optimal rule in each world. When transport costs are very high, each location operates in autarky. There are no spillovers to unflooded locations and no heterogeneity in loss among the flooded. The optimal policy coincides with a flat transfer to the physically exposed. When transport costs are very low, the productivity loss diffuses uniformly across the economy. The optimal policy still varies transfers across locations, but with little heterogeneity in loss there is little to gain over a flat transfer. The gains from optimal targeting are largest in between, where losses neither stay local nor diffuse away – they peak at around 12%. This is the empirically relevant case, since real economies sit between the two extremes. Figure 6 shows

the three cases. Under low integration the loss stays local and every flooded location is hit alike; under high integration it is close to uniform; in the calibrated middle it concentrates on the interior locations far from dry land, which receive the largest transfers.

Figure 6: Welfare loss and the optimal aid allocation under alternative trade costs



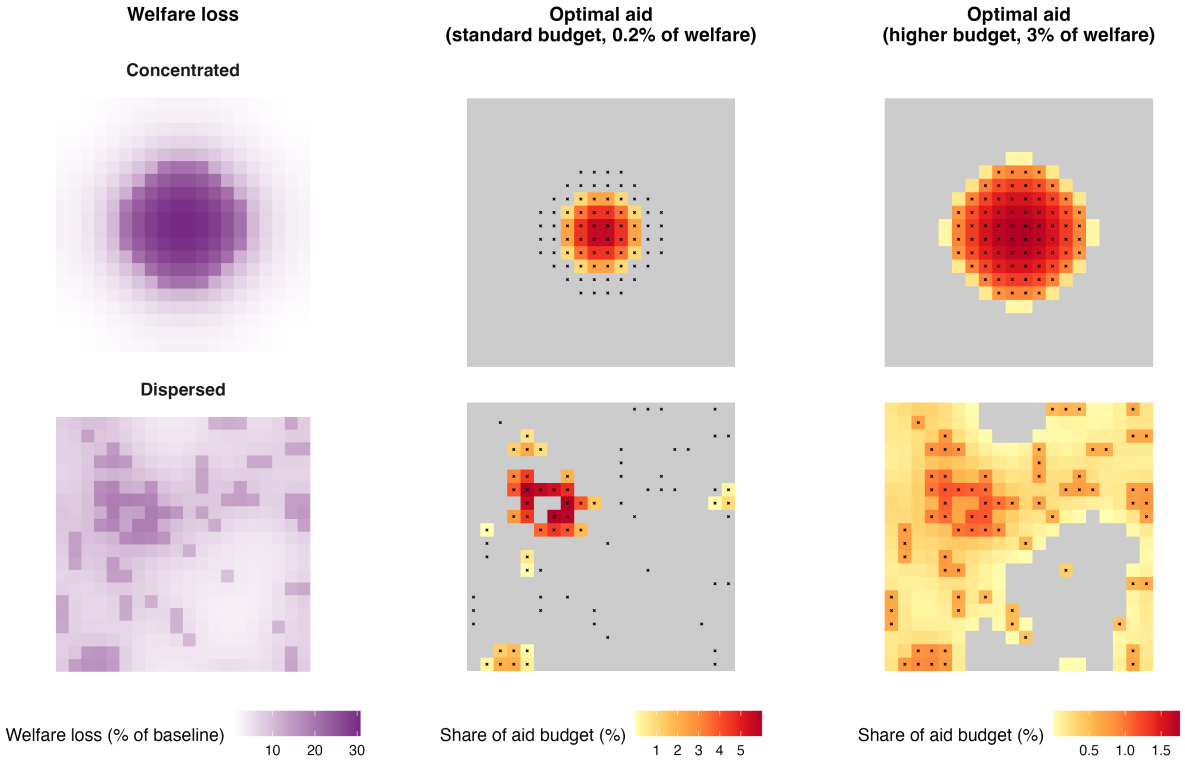
Notes: Columns are three levels of economic integration, from least integrated (left) to most integrated (right). Top row: worker welfare loss (% of baseline). Bottom row: the continuous-optimal aid allocation (each location's share of the aid budget, on a common scale across panels; the budget is the same in every world); crosses mark flooded locations.

Spatial concentration

Holding size fixed, the spatial pattern of a flood changes who bears the loss and where aid should go. A dispersed flood leaves most flooded locations with unflooded neighbours to commute to and trade with, so the loss spills onto those neighbours – 65% of it falls outside the physical flood zone, against 40% under a concentrated flood of the same size (Figure 7, left column). A concentrated flood does the reverse. Its centre is far from dry land, while its edge can still reach the unflooded locations nearby. The result is that not only is more of the loss born by the flooded, but it is also more unequally distributed amongst them - the worst-hit interior location loses 31% of baseline welfare, against 19% in the dispersed flood.

This steep gradient is what optimal transfers are able to target. Under a concentrated flood the optimal policy gives most resources to the centre and little to the edge (Figure 7, centre column), and the welfare gain over a flat transfer is roughly 2.4 times larger than for a dispersed flood of the same size. A dispersed flood spreads the loss more evenly, leaving less for targeting to do within the zone. At a larger budget the optimal policy begins to reach beyond the flood zone, and does so far more under the dispersed flood, which has lowered welfare in the unflooded locations nearby (Figure 7, right column).

Figure 7: Flood incidence and optimal aid, by flood dispersion



Notes: Two floods of the same size ($\phi = 0.19$): concentrated (top row) and dispersed (bottom row). Left column: worker welfare loss (% of baseline). Centre and right columns: the continuous-optimal aid allocation (each location's share of the aid budget), at the standard budget (0.2% of pre-flood worker welfare) and at a higher budget (3%); crosses mark flooded locations.

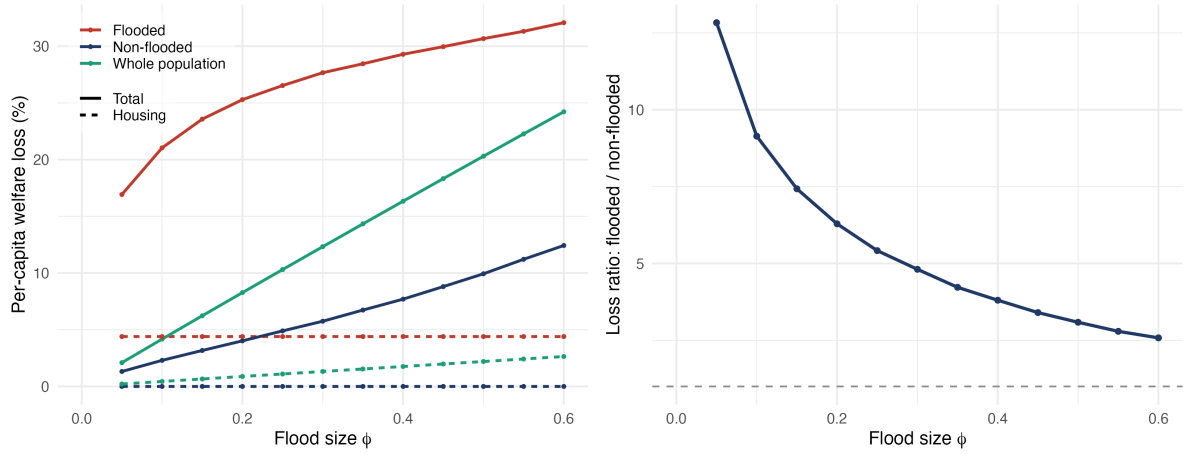
Size of the flood

The larger the flood, the smaller the relative gap between what a flooded and an unflooded location loses. The directly flooded almost always lose more, for two reasons. They suffer housing damage, which is local and unshared, and flooding is spatially correlated, so a flooded location is more likely

to be far from dry land, so it loses more market access. At a representative flood a flooded location loses 25% of baseline welfare, three points of it housing and the rest productivity, against 4% for an unflooded location, all of it productivity (Figure 8, left).

The housing part of this gap is a fixed local loss that does not grow with the flood. As the flood spreads, the productivity loss spreads more evenly across flooded and unflooded locations, so the gap between them shrinks – the ratio of flooded to unflooded loss falls from 9 at a small flood to 3 at a large one (Figure 8, right). Physical exposure therefore reveals less about who bears the loss as the flood grows.

Figure 8: Flood size and the distribution of losses between flooded and unflooded households



Notes: Both panels sweep over flood size ϕ (concentrated discs, calibrated frictions). Left: per-capita worker welfare loss for directly flooded, unflooded, and whole-population locations, with the housing own-cell component (dashed) separated from the total (solid). Right: the ratio of average loss, flooded to unflooded, as ϕ rises.

Two events

The three factors combine differently across disasters. Hurricane Katrina in 2005 was geographically concentrated and flooded a small share of the country. In our terms it is a small, concentrated flood – losses are unequal among the affected, so targeting within the flooded area can help, but the spillovers onto unaffected areas are limited, and evidence finds little effect on the labour markets of neighbouring locations (Deryugina et al., 2018). There is little to gain from extending aid beyond the directly exposed. Pakistan's 2022 floods are the opposite – very large, with roughly a third of the country and a fifth of Sindh under water, and reasonably dispersed. Here the model predicts the

widespread spillovers we find, and little gap between directly and indirectly exposed households, so aid targeted on physical exposure alone misses much of where the loss actually falls.

8 Conclusion

This paper asks how the economic costs of a flood depend not only on whether a household is itself inundated, but on how much of the surrounding area floods with it. We combine a panel of 5,100 rural households in Sindh with a recentered instrumental-variables design that separately identifies the impact of direct and indirect flood exposure. We find that physical exposure is a poor guide to economic incidence. Direct exposure causes physical damage, but on its own it leaves employment, prices, credit and consumption largely unmoved. By contrast, indirect exposure lowers employment, raises food prices, and tightens credit – a household surrounded by flooding is less likely to be in work, faces higher food prices, and finds credit harder to come by.

Aid, however, did not follow these losses. Directly exposed households were more likely to receive it, while indirectly exposed households received less. The households hit on both margins bore the largest costs, yet received less than those with dry markets to fall back on. Relief organised around observable damage to homes and infrastructure therefore misses much of the incidence of a large flood, and underserves those who bear most of it. In a model of these spillovers, optimal relief directs more to the worst-hit, typically the interior of the flooded zone, far from dry land.

Informal coping works best against localised shocks, and breaks down when the shock is large enough to engulf the markets and neighbours a household would otherwise turn to. As climate change makes large floods more frequent, and a single event covers more of a country at once, the indirect costs of widespread flooding – and the gap left by relief that ignores them – are likely to become an increasingly central concern for disaster policy.

References

Aiken, Emily, Suzanne Bellue, Dean Karlan, Christopher Udry, and Joshua E. Blumenstock, "Machine Learning and Phone Data Can Improve Targeting of Humanitarian Aid," *Nature*, 2022, *603* (7903), 864–870. [Return to page 6]

Alatas, Vivi, Abhijit Banerjee, Rema Hanna, Benjamin A. Olken, and Julia Tobias, "Targeting the Poor: Evidence from a Field Experiment in Indonesia," *American Economic Review*, 2012, *102* (4), 1206–1240. [Return to page 6]

_ , Ririn Purnamasari, Matthew Wai-Poi, Abhijit Banerjee, Benjamin A. Olken, and Rema Hanna, "Self-Targeting: Evidence from a Field Experiment in Indonesia," *Journal of Political Economy*, 2016, *124* (2), 371–427. [Return to page 6]

Allen, Treb and David Atkin, "Volatility and the Gains from Trade," *Econometrica*, 2022, *90* (5), 2053–2092. [Return to page 6]

Baily, Martin Neil, "Some Aspects of Optimal Unemployment Insurance," *Journal of Public Economics*, 1978, *10* (3), 379–402. [Return to page 36]

Balboni, Clare, Johannes Boehm, and Mazhar Waseem, "Firm Adaptation in Production Networks: Evidence from Extreme Weather Events in Pakistan," *American Economic Review*, 2025. Forthcoming. [Return to page 6]

Banerjee, Abhijit, Rema Hanna, Benjamin A. Olken, and Diana Sverdlin Lisker, "Social Protection in the Developing World," *Journal of Economic Literature*, 2024, *62* (4), 1349–1421. [Return to page 6]

Banerjee, Abhirup, Shraddha Gupta, Pranava Priyanshu, Ankan Kar, Ruby Saha, Tanujit Chakraborty, Dibakar Ghosh, Jürgen Kurths, and Chittaranjan Hens, "Recent changes in spatiotemporal patterns of heat extremes in South Asia," *npj Climate and Atmospheric Science*, 2025, *8* (293). [Return to page 8]

Barrot, Jean-Noël and Julien Sauvagnat, "Input Specificity and the Propagation of Idiosyncratic Shocks in Production Networks," *The Quarterly Journal of Economics*, 2016, *131* (3), 1543–1592. [Return to page 6]

Baseler, Travis and Jakob Hennig, "Disastrous Displacement: The Long-Run Impacts of Landslides," *American Economic Journal: Applied Economics*, 2023. Forthcoming; World Bank Policy Research Working Paper 10535. [Return to page 6]

Boehm, Christoph E., Aaron Flaaen, and Nitya Pandalai-Nayar, "Input Linkages and the Transmission of Shocks: Firm-Level Evidence from the 2011 Tōhoku Earthquake," *The Review of Economics and Statistics*, 2019, *101* (1), 60–75. [Return to page 6]

Borusyak, Kirill and Peter Hull, "Nonrandom Exposure to Exogenous Shocks," *Econometrica*, 2023, *91* (6), 2155–2185. [Referenced on pages 2 and 11]

_ **and** _ , "Optimal Formula Instruments," Working Paper 33594, National Bureau of Economic Research 2025. [Return to page 14]

Boustan, Leah Platt, Matthew E. Kahn, Paul W. Rhode, and Maria Lucia Yanguas, "The Effect of Natural Disasters on Economic Activity in US Counties: A Century of Data," *Journal of Urban Economics*, 2020, *118*, 103257. [Return to page 6]

Breza, Emily, Supreet Kaur, and Yogita Shamdasani, "Labor Rationing," *American Economic Review*, 2021, *111* (10), 3184–3224. [Return to page 3]

Brunner, Manuela I. and Corentin Chartier-Rescan, "Drought Spatial Extent and Dependence Increase During Drought Propagation From the Atmosphere to the Hydrosphere," *Geophysical Research Letters*, 2024, *51* (e2023GL107918). [Return to page 8]

Burgess, Robin and Dave Donaldson, "Can Openness Mitigate the Effects of Weather Shocks? Evidence from India's Famine Era," *American Economic Review*, 2010, *100* (2), 449–453. [Return to page 6]

Carvalho, Vasco M., Makoto Nirei, Yukiko U. Saito, and Alireza Tahbaz-Salehi, "Supply Chain Disruptions: Evidence from the Great East Japan Earthquake," *The Quarterly Journal of Economics*, 2021, *136* (2), 1255–1321. [Return to page 6]

Castro-Vincenzi, Juan Manuel, "Climate Hazards and Resilience in the Global Car Industry," Working Paper, Princeton University 2024. [Return to page 6]

Chamberlain, Gary, "Asymptotic efficiency in estimation with conditional moment restrictions," *Journal of Econometrics*, 1987, *34* (3), 305–334. [Return to page 14]

Chernozhukov, Victor, Denis Chetverikov, Mert Demirer, Esther Duflo, Christian Hansen, Whitney Newey, and James Robins, "Double/debiased machine learning for treatment and structural parameters," *The Econometrics Journal*, 2018, *21* (1), C1–C68. [Return to page 14]

Chetty, Raj, "A General Formula for the Optimal Level of Social Insurance," *Journal of Public Economics*, 2006, *90* (10–11), 1879–1901. [Return to page 36]

Deryugina, Tatyana, Laura Kawano, and Steven Levitt, "The Economic Impact of Hurricane Katrina on Its Victims: Evidence from Individual Tax Returns," *American Economic Journal: Applied Economics*, 2018, *10* (2), 202–233. [Referenced on pages 6 and 45]

Fu, Chao and Jesse Gregory, "Estimation of an Equilibrium Model with Externalities: Post-Disaster Neighborhood Rebuilding," *Econometrica*, 2019, *87* (2), 387–421. [Return to page 6]

Funk, Chris, Pete Peterson, Martin Landsfeld, Diego Pedreros, James Verdin, Shraddhanand Shukla, Gregory Husak, James Rowland, Laura Harrison, Andrew Hoell, and Joel Michaelsen, "The climate hazards infrared precipitation with stations—a new environmental record for monitoring extremes," *Scientific Data*, 2015, *2* (1), 150066. [Return to page 7]

Gallagher, Justin and Daniel Hartley, "Household Finance after a Natural Disaster: The Case of Hurricane Katrina," *American Economic Journal: Economic Policy*, 2017, *9* (3), 199–228. [Return to page 6]

Gignoux, Jérémie and Marta Menéndez, "Benefit in the Wake of Disaster: Long-Run Effects of Earthquakes on Welfare in Rural Indonesia," *Journal of Development Economics*, 2016, *118*, 26–44. [Return to page 6]

Gordon, Matthew, Yukiko Hashida, and Eli P. Fenichel, "Targeting Disaster Aid: A Structural Evaluation of a Large Earthquake Reconstruction Program," Technical Report, SSRN Working Paper 4862074 2025. [Return to page 6]

Groen, Jeffrey A., Mark J. Kutzbach, and Anne E. Polivka, "Storms and Jobs: The Effect of Hurricanes on Individuals' Employment and Earnings over the Long Term," *Journal of Labor Economics*, 2020, *38* (3), 653–685. [Return to page 6]

Hanna, Rema and Benjamin A. Olken, "Universal Basic Incomes versus Targeted Transfers: Anti-Poverty Programs in Developing Countries," *Journal of Economic Perspectives*, 2018, *32* (4), 201–226. [Return to page 6]

Huber, Kilian, "Estimating General Equilibrium Spillovers of Large-Scale Shocks," *The Review of Financial Studies*, 2023, *36* (4), 1548–1584. [Return to page 18]

Kemter, Matthias, Bruno Merz, Norbert Marwan, Sergiy Vorogushyn, and Günter Blöschl, "Joint Trends in Flood Magnitudes and Spatial Extents Across Europe," *Geophysical Research Letters*, 2020, *47* (e2020GL087464). [Return to page 8]

Kirchberger, Martina, "Natural Disasters and Labor Markets," *Journal of Development Economics*, 2017, *125*, 40–58. [Return to page 6]

Kocornik-Mina, Adriana, Thomas K. J. McDermott, Guy Michaels, and Ferdinand Rauch, "Flooded Cities," *American Economic Journal: Applied Economics*, 2020, *12* (2), 35–66. [Return to page 6]

McIntosh, Molly Fifer, "Measuring the Labor Market Impacts of Hurricane Katrina Migration: Evidence from Houston, Texas," *American Economic Review*, 2008, *98* (2), 54–57. [Return to page 6]

Nigro, Joseph, Daniel Slayback, Frederick Policelli, and G. Robert Brakenridge, "NASA/DFO MODIS Near Real-Time (NRT) Global Flood Mapping Product Evaluation of Flood and Permanent Water Detection," Technical Report, NASA Goddard Space Flight Center, Greenbelt, MD 2014. [Return to page 7]

Patel, Dev, "Floods," 2024. Working paper. [Referenced on pages **6** and **10**]

Pople, Ashley, Ruth Hill, Stefan Dercon, and Ben Brunckhorst, "Importance of Being Early? Anticipatory Cash Transfers for Flood-Affected Households," *The World Bank Economic Review*, 2026. Advance access. [Return to page **6**]

Qin, Lianjie, Laiyin Zhu, Baoyin Liu, Zixuan Li, Yugang Tian, Gordon Mitchell, Shifei Shen, Wei Xu, and Jianguo Chen, "Global expansion of tropical cyclone precipitation footprint," *Nature Communications*, 2024, *15* (4824). [Return to page **8**]

Roxy, M. K., Subimal Ghosh, Amey Pathak, R. Athulya, Milind Mujumdar, Raghu Murtugudde, Pascal Terray, and M. Rajeevan, "A threefold rise in widespread extreme rain events over central India," *Nature Communications*, 2017, *8* (708). [Return to page **8**]

Simonovska, Ina and Michael E. Waugh, "The Elasticity of Trade: Estimates and Evidence," *Journal of International Economics*, 2014, *92* (1), 34–50. [Return to page **34**]

Skinner, Christopher B., Danielle Touma, Mathew Barlow, Deepti Singh, and Troy King, "The spatial extent of heat waves has changed over the past four decades," *Communications Earth & Environment*, 2025, *6* (662). [Return to page **8**]

Sotelo, Sebastian, "Domestic Trade Frictions and Agriculture," *Journal of Political Economy*, 2020, *128* (7), 2690–2738. [Return to page **34**]

Tarquinio, Lisa, "The Politics of Drought Relief: Evidence from Southern India," Technical Report, SSRN Working Paper 4163045 2023. [Return to page **6**]

Tellman, B., J. A. Sullivan, C. Kuhn, A. J. Kettner, C. S. Doyle, G. R. Brakenridge, T. A. Erickson, and D. A. Slayback, "Satellite imaging reveals increased proportion of population exposed to floods," *Nature*, 2021, *596* (7870), 80–86. [Return to page **7**]

World Weather Attribution, "Climate change likely increased extreme monsoon rainfall, flooding highly vulnerable communities in Pakistan," https://www.worldweatherattribution.org/climate-change-likely-increased-extreme-monsoon-rainfall-flooding-highly-vulnerable-communities-in-pakist September 2022. Scientific attribution statement released 14 September 2022. Attributes a 75%

increase in 60-day Sindh/Balochistan rainfall intensity to anthropogenic warming. [Return to page 7]

A Appendix figures

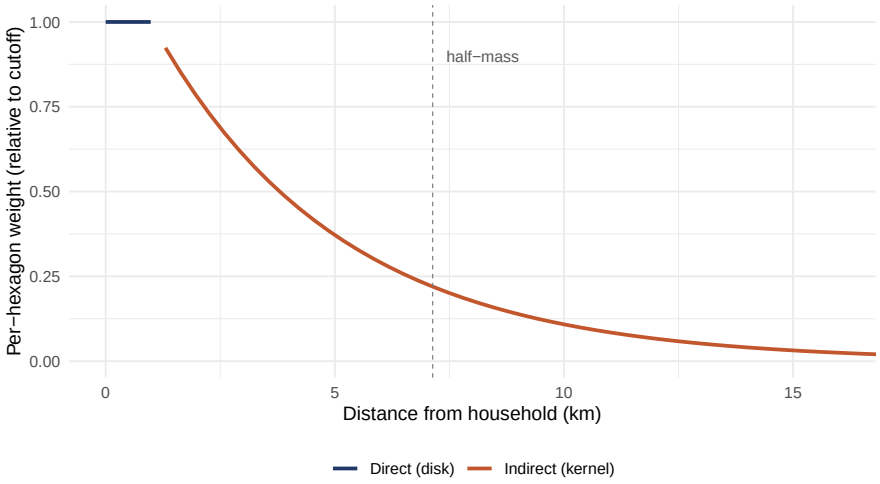


Figure 9: Exposure weights by distance