

The Economic Incidence of Natural Disasters

Amen Jalal*

Pol Simpson[‡]

October 28, 2026

This draft is updated frequently. Click [here](#) for the latest version.

Abstract

Policymakers and researchers often view the costs of natural disasters as local to the areas they physically strike. We show that due to spatial linkages across markets, the economic incidence of a disaster extends well beyond its physical footprint. We study Pakistan's large-scale 2022 floods, combining original panel surveys of 5,100 rural households with satellite flood maps and a recentered instrument based on a hydrological model. We find that moving a household from the 10th to the 90th percentile of flooding within 15km of its home lowers its real expenditure by 10%, of which nearly two-thirds is due to flooding more than 5km away from home. The same change in flooding reduces the share of men in wage employment by 9 percentage points (pp), raises food prices by 8% and asset sales by 3 pp, and lowers credit access by 5 pp. These effects are driven almost entirely by flooding more than 5km away. We quantify the economic incidence of flooding in general equilibrium by building a spatial model of trade and commuting. Although two in five locations experienced almost no flooding, the model implies that welfare fell everywhere, and that 73% of the average loss was caused by flooding elsewhere. Moreover, targeting aid on predicted losses rather than physical exposure achieves the same welfare at 18% lower cost. Thus, ignoring the divergence between economic incidence and physical exposure leads us to underestimate disaster costs and misallocate aid, particularly for larger disasters and more integrated economies.

*Harvard University; amenjalal@fas.harvard.edu.

[†]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 our advisors Oriana Bandiera, Gharad Bryan, Robin Burgess and Daniel Sturm for invaluable support. For very helpful comments and feedback, we also thank David Atkin, Clare Balboni, Maarten De Ridder, Wouter Den Haan, Matthew Gordon, Allison Green, Allan Hsiao, Felix Iglhaut, Rohini Pande, Dev Patel, Steve Pischke, Emiliano Rinaldi, Kartik Srivastava and seminar participants at the 2026 CSQIEP Mentoring Conference, LSE Economics of Environment and Energy (EEE) hour, LSE Environment Camp, LSE Environment Week and NBER Summer Institute. Zahra Tajammal, Isra Imtiaz and Maheen Memon provided outstanding research assistance.

1 INTRODUCTION

Governments, aid agencies, and researchers often view the impacts of natural disasters as local, assuming that those directly hit bear most of the costs. Yet trade can transmit losses far beyond the locations a disaster physically strikes. Flooded households, for instance, may buy goods and find work in drier areas. This provides them with partial insurance against their own losses but shifts some of the costs onto residents elsewhere. If these spillovers are large, ignoring them will lead us to understate the impacts of natural disasters and misallocate aid.

This paper estimates the economic incidence of Pakistan’s 2022 floods in general equilibrium. These floods were the largest in the country’s 75-year history and were induced by unprecedented rainfall. In Sindh, the hardest-hit province, summer rain was 267% above its long-run average. Nationwide, the floods inundated 7.3% of Pakistan’s land and directly affected 14% of its population. Globally, these floods rank in the top 5% by share of national land inundated and the top 1% by share of population affected, compared to other major floods in the past quarter century.

In this setting, we first show that, empirically, the economic incidence of flooding differs substantially from its physical footprint. To do so, we draw a random sample of 5,100 low-income rural households from a pre-flood census covering six districts in Sindh. We then locate, follow, and survey these households over the two years following the floods. We combine our survey data with satellite flood maps and a recentered instrumental variables (IV) design ([Borusyak and Hull, 2023](#)). The IV exploits deviations between actual 2022 flooding and expected flooding at each location, where the latter is simulated using a hydrological model and satellite data on historical rainfall.

We then build a spatial model to quantify the general-equilibrium effects of flooding. To assess how broadly our findings generalise, we also simulate counterfactual floods of different magnitudes and under different levels of market integration. Finally, we use the model to study how spillovers shape the efficient allocation of disaster aid, allowing for the possibility that aid itself generates general-equilibrium effects.

Our reduced-form estimates show that a household’s expenditure loss depends more on flooding in the surrounding area than on flooding at its own location. We compare households at the 90th and 10th percentiles of flood exposure, and measure exposure as the extent of flooding in a given radius.

When the radius is just 1 km around a home, real expenditure is 6.8% lower for highly exposed households ($p = 0.05$). Widening the radius raises the estimated loss to 9.4% at 10 km ($p = 0.02$) and 10.1% at 15 km ($p = 0.01$). We then decompose the total effect of flooding within 15 km into the contribution of nearer and more distant flooding, by including separate regressors for flooding within a cutoff distance of x km and flooding beyond that cutoff, up to 15 km. This exercise shows that flooding within 1 km of home accounts for only about 11% of the expenditure decline, while flooding more than 5 km away accounts for 64%.

Why does flooding several kilometres from a household affect its consumption? We find that the answer lies in linkages through labour, product, and credit markets across space. Under autarky, only flooding at a household's own location would matter. In practice, access to nearby locations helps households cope with local losses, unless they are also flooded.

Commuting illustrates how these connections work. A 10th–90th percentile increase in flooding within 1 km of a house raises by 6 pp the likelihood that a household member commutes more than 1 km to work. This suggests that households respond to local flooding by seeking work farther afield. But the effect shrinks as the exposure radius widens, reaching roughly zero at 15 km. The distance decomposition helps explain why. Holding flooding in other distance bands constant, flooding within 5 km raises commuting by about 4 pp. By contrast, flooding 5–15 km away lowers it enough to offset that increase. In other words, the commuting response to localised flooding is offset when flooding also affects the surrounding area.

Employment shows a similar pattern. A 10th–90th percentile increase in flooding within 15 km reduces men's wage employment by 9 pp, with virtually all of the decline due to flooding beyond 1 km. Over the same radius, the food-price index rises by 8%, the share of households with access to loans falls by 5 pp, and asset sales rise by about 3 pp. Flooding beyond 1 km accounts for nearly all of the price and loan effects and most of the increase in asset sales. Together, these findings suggest that spatial links can help households adjust to localised flooding, while spreading its costs across the region.

A potential concern with our analysis is that distant flooding may proxy for local flooding when household locations are measured with error, inflating our estimates of spillovers. We test for this using housing damage as a placebo outcome. Our measure comes from a geo-tagged census of

properties that the Government of Sindh conducted after the floods to assess damage and guide reconstruction. Housing damage is a useful placebo because a home cannot be damaged by floodwaters far away, so any estimated effect of distant flooding on damage would suggest measurement error. We find no such effect: flooding within 1 km raises the probability of damage by 6 pp, while flooding farther away has no detectable effect.

To quantify the total effect of the flood on welfare in each location, we use a spatial model in which locations are linked through costly trade and commuting, based on [Eaton and Kortum \(2002\)](#) and [Ahlfeldt et al. \(2015\)](#). There are two sectors of production, agriculture, which uses labour and land, and non-agriculture, which only uses labour. Consistent with our data, workers live in a fixed location and choose where and in which sector to work. Worker welfare depends on access to well paying jobs, cheap goods and their housing endowment. Flooding lowers agricultural productivity and destroys housing in affected locations. As workers adjust where they buy and work, prices and wages change throughout the economy. The extent of these spillovers depends on pre-flood economic integration and on how readily buyers and workers can substitute across locations.

We discipline these links using data from our setting, starting with bilateral travel times computed from the road network. We use internal trade data to estimate how trade flows vary with travel time, then separately identify the trade elasticity by regressing post-flood prices on flood-induced changes in market access using a recentered instrument. We estimate how commute flows vary with travel time using data from our surveys, and calibrate the commuting elasticity to match the observed dispersion of wages.

We quantify the model for Sindh, which we divide into approximately 14,000 hexagons, each with a radius of 1.2 km. Using data on agricultural revenues, population, and agricultural employment, we recover pre-flood productivity in each location. We then combine the flood map with the estimated effects of flooding on agricultural output and housing damage to compute post-flood fundamentals, and solve for the resulting equilibrium. This approach allows us to map flooding in any set of locations into welfare losses across the entire province.

The model results show that the economic impacts of the 2022 floods extended well beyond the locations they inundated. Two in five locations experienced almost no flooding, yet welfare fell in every location, by at least 1.6% and by 12.2% on average. For the average location, 73% of the

welfare loss came from flooding elsewhere. This share was 99% in the least-flooded decile and 32% in the most-flooded decile.

Counterfactual floods show how the distribution of losses changes with a disaster's scale and spatial pattern. When 5% of locations are flooded, the welfare loss per resident in a flooded location is 3.8 times that in an unflooded location. When 60% are flooded, this ratio falls to 2.1, as losses in unflooded locations nearly triple. Dispersing flooding across space raises the share of total losses borne by unflooded locations from 24% to 30%.

We also vary market integration to assess how trade and commuting links affect the reach of these losses. These counterfactuals clarify why comparisons between more- and less-exposed households measure relative effects: when flooding affects outcomes in both groups, the comparison nets out part of the total impact. This gap grows with market integration and with flood size over the range we study, attenuating the estimated flood impacts.

The distinction between economic incidence and physical exposure matters for the allocation of disaster aid. A rank-rank regression of model-predicted welfare losses on own flooding has an R^2 of 0.71, so physical exposure predicts losses but does not fully capture them. A rule that transfers aid to every flooded location and one that targets an equal number of locations selected by predicted welfare loss disagree for one in six locations. Those missed by the exposure-based rule tend to be less-flooded locations surrounded by extensive flooding, for whom flooding elsewhere causes 92% of their loss.

To assess the potential gains from targeting aid according to economic incidence we introduce a planner who allocates a fixed budget to minimise a convex function of flood-induced welfare losses. Compared to targeting transfers proportional to a location's own flooding, targeting on model-predicted losses could achieve the same welfare at 82% of the cost. These savings depend on the planner's ability to target losses directly, and fall to 3% if transfers can be conditioned only on own flooding. These calculations assume that aid itself has no general equilibrium effects, which we then relax, allowing aid that restores productive capacity to change wages and prices through the same spatial links that transmit flood losses.

This paper builds on a literature that shows that trade reallocates risk across space and that economic integration changes the exposure of locations to local shocks ([Newbery and Stiglitz, 1984](#);

Burgess and Donaldson, 2010; Allen and Atkin, 2022; Brooks and Donovan, 2020). For example, Allen and Atkin (2022) show that highway expansion in India made local prices less responsive to local yields and more responsive to yields elsewhere, and Brooks and Donovan (2020) show that bridges eliminate the response of labour earnings to local flooding in Nicaraguan villages. We also draw on work on the informal mechanisms households use to insure against shocks, whether individually, through migration (Morten, 2019), or collectively, through community and network risk-sharing (Townsend, 1994; Munshi and Rosenzweig, 2016).

Building on these insights, we contribute principally to the literature on the impacts of natural disasters. This literature has documented large direct effects of disasters on households (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), and has shown that firms can be affected through supply chains even when they are not physically exposed (Barrot and Sauvagnat, 2016; Carvalho et al., 2021; Balboni et al., 2025; Boehm et al., 2019; Castro-Vincenzi, 2024). These literatures study direct and indirect exposure separately. We estimate their effects on households jointly, and quantify the overall incidence of a disaster in general equilibrium. In doing so, we show that indirect exposure accounts for most of the welfare loss from a large flood. This implies that research designs focused on direct physical exposure may substantially underestimate the total impact of natural disasters.

We also contribute to the literature on the targeting of social programmes, which studies how best to allocate resources when the planner faces information constraints (Ravallion, 1993; Elbers et al., 2007; Alatas et al., 2012; Hanna and Olken, 2018; Aiken et al., 2022). In our context, the planner uses transfers to offset the impact of floods. The losses from a flood are causal effects, determined in general equilibrium, and are not directly observable. Our problem therefore relates closely to recent work on targeting transfers by their estimated treatment effects (Kitagawa and Tetenov, 2018; Athey and Wager, 2021; Haushofer et al., 2025), and to the literature on optimal spatial transfers, which uses a model-based approach to measure the general-equilibrium impact of transfers on economic geography (Fajgelbaum and Gaubert, 2020; Donald et al., 2026). We develop a feasible approach to targeting aid on these losses, using a model quantified with widely available data, which can also account for the general-equilibrium effects of the transfers themselves. In the context of disasters,

[Gordon et al. \(2025\)](#) find that property damage is a poor basis for targeting reconstruction aid within affected villages, due to household-level heterogeneity in access to self-insurance. We show that physical exposure is also an imperfect basis for targeting, for a different reason, since it misses losses that propagate across space through markets.

The rest of the paper proceeds as follows. Section 2 describes the 2022 floods and Section 3 the data. Section 4 sets out the research design and Section 5 presents the reduced-form results. Section 6 develops, estimates and quantifies the spatial model. Section 7 evaluates the incidence of the 2022 floods and of counterfactual floods, and Section 8 studies the targeting of disaster assistance. Section 9 concludes.

2 PAKISTAN’S 2022 FLOODS

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

Heavy rain is not unusual in Pakistan, which experiences a monsoon every summer. However, in 2022 the cumulative July–August rainfall was the greatest it had been 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 267% above the long-run average, and 44% above its own previous record (2003). It is very likely that climate change contributed to this unprecedented event. An 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 rains caused widespread flooding. According to MODIS satellite imagery ([Nigro et al., 2014](#)) 7.3% of normally dry land in Pakistan was inundated at some point between July and December 2022. In Sindh the figure is 20.4%, and in the six districts where we work it reached 41%. Because Sindh is low-lying and flat, the water had nowhere to drain, and large areas of the province stayed inundated for months whilst the waters slowly evaporated.

The floods were extreme by global standards, not only Pakistan’s. The Global Flood database of [Tellman et al. \(2021\)](#) collates information on the 906 largest floods worldwide since 2000, taking as its starting point the Dartmouth Flood Observatory’s log of floods reported in news and government

sources. For each flood, it uses satellite imagery to record the share of national land and population that was inundated. Across these floods, a flood at the 90th percentile inundated 2.8% of national land and directly affected 1.7% of national population. Pakistan’s 2022 floods covered 7.3% of the country’s land and directly affected 14% of its people, placing this event at the 95th and 99th percentiles respectively¹.

This large-scale event is a compelling setting for our research question. In part this is a matter of power since 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 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

3 DATA

To estimate how household outcomes respond to flooding at different distances, we need geolocated data on consumption, employment and prices. Data of this kind is hard to come by in Pakistan. Administrative data has limited coverage or is not regularly updated. The government suspended its household surveys for census preparations, and those it does field ask few of the questions we need. Satellite data does not record prices, consumption or incomes.

We therefore bring together data from a wide range of sources with granular geographic information, down to points or hexagons a kilometre across. Our principal source is an original, geocoded panel of rural households, which we collected to measure household outcomes at the level our question requires. To measure physical damage, we combine official microdata on housing damage with satellite building footprints. For agricultural output and revenues, we classify plots by crop type using satellite data and machine learning, proxy output with NDVI, calibrate to district-level agricultural surveys, and convert to revenues using prices from major markets. We measure flooding from multiple satellite sources.

¹It is clear from [Patel \(2024\)](#) that many floods, particularly smaller ones, are omitted from the DFO. This suggests we may be underestimating how extreme Pakistan’s 2022 floods were in the global context.

²The 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.

This section also describes the topographical and rainfall data underlying our identification strategy, and the internal trade, transport network, population and labour force data we use to quantify the model. Table 3.1 summarises the sources.

Table 3.1: Data sources

Variable	Source	Years	Resolution
<i>Household survey</i>			
Household outcomes	Original panel	2017, 2023, 2024	Household (GPS)
<i>Physical damage and agriculture</i>			
Building damage rate	Damage census; VIDA	2022–23	Hex ~0.9km (res. 8)
Agricultural revenues	Sentinel-1/2; MODIS; SBOS; SPI	2018–2024	Hex ~2.4km (res. 7)
Agricultural employment share	Damage census; PLFS	2020–23	Hex ~0.9km (res. 8)
<i>Flooding</i>			
Flood index	MODIS; Sentinel-1	Jul–Dec 2022	Hex ~0.9km (res. 8)
<i>Hydrological model inputs</i>			
Rainfall	IMERG; ECMWF forecasts	1998–2025	0.1°–0.4°, daily
Topography	FABDEM	Static	30m
<i>Model inputs</i>			
Trade flows	Market arrival records	2018–2026	Market, daily
Road network	OpenStreetMap	Static	Road segment
Population	WorldPop	2020	100m
Housing stock	GHSL built-up surface	2020	100m

3.1 Household survey data

We collect panel data on 5,100 poor, rural households in Sindh, first observed in 2017 and surveyed twice after the floods, in summer 2023 and 2024.

We draw the sample from a census of all low-income households in six districts of the province (Qambar Shahdadt, Larkana, Matiari, Tando Muhammad Khan, Tando Allahyar and Sujawal). Whilst restricting attention to the rural poor limits our ability to comment on the experience of wealthier or urban households, its key benefit is that we can compare very similar households that happened to face very different flooding.

A local NGO collected the data to determine eligibility for development and poverty alleviation programmes. It 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. In addition to treating this data as a sampling frame, we

extract baseline covariates from it, which we use as controls to improve precision and to run balance checks.

We sample to span the range of flood exposure, choosing survey locations within each district with high and low predicted flooding, and then randomly drawing households from them. We survey households in two waves, running July–September of 2023 and 2024. With the assistance of local community members, we locate 92% of households in 2023 or 2024; of those, almost all are still living in their 2017 village.³ Attrition is low and does not vary with 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 4,683 households surveyed in both waves.

The survey is extensive, so we highlight the modules that give us our main outcomes. Total expenditure and consumer prices come from a consumption module, which records total spending alongside the quantity and value consumed of meat, dairy, fruit and legumes. Employment status, occupation and wages come from a household roster, which covers demographics, education and labour for every member. Receipt of flood relief comes from a relief module, which records the type and provider of assistance alongside self-reported flood exposure and evacuation decisions. We also collect commuting data, which we use to estimate and calibrate the commuting side of the model. In the first wave we record each worker’s travel time to work. In the second we additionally record their specific place of work, monetary commuting costs and commute mode.

3.2 Administrative damage data

To measure physical damage to housing we use the Government of Sindh’s in-person census of damaged residential property, which it conducted after the floods to administer reconstruction grants, and which records 2.1 million properties. To build a damage rate, we aggregate the census to hexagons of radius ~460m (resolution 8 of Uber’s H3 system) and divide by the buildings the VIDA footprint layer counts there.

³We 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 census also records whether the main earner of the household works in agriculture, which we combine with the Pakistan Labour Force Survey (PLFS) to measure the agricultural share of employment in each hexagon, which the model needs. For rural hexagons we average the microdata by hexagon, then shift every cell in a district by a common amount so that the district average matches the PLFS rural share. For urban hexagons we use the PLFS urban share for that district.

3.3 Agricultural revenues

We measure agricultural revenue in each hexagon of radius $\sim 1.2\text{km}$ (resolution 7), by season, from 2018 to 2024, by combining satellite data with official production statistics. We first use the satellite data to establish which crops grow where, and to build a local proxy for output. We identify the cropland in each hexagon from ESA WorldCover, a 10m land-cover map. We then classify that cropland by crop, training a random forest on plot-level data from the ADB Ground Truthing Survey to predict rice, cotton, sugarcane and maize from how a plot looks to radar and optical satellites across the season. Held-out accuracy is 91%. Multiplying cropped area by NDVI, a satellite index of greenness, in September and February of each year gives us a proxy for Kharif and Rabi output. Since this proxy is measured in greenness-weighted cropped area rather than tonnes, we rescale it so that the total for each district, crop and year matches the tonnage the Sindh Bureau of Statistics reports. We then value the tonnes at crop prices from the Sensitive Price Indicator. This gives us agricultural revenue for each hexagon in each Kharif and Rabi season.

3.4 Flood data and flood index construction

We measure flood extent and duration using 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 each location less frequently.⁴ We measure depth by combining the Sentinel-1 water maps with the FLEXTTH algorithm (Betterle and Salamon, 2024). Intuitively, at the edge of a flood the water surface sits at ground level, so the flood boundary reveals the height of the flood water above sea level, and the depth at any interior point is that water level minus the ground elevation.

⁴For an in-depth discussion of the challenges of measuring floods from space, see Patel (2024).

We build a single index of flood intensity covering three dimensions along which floods differ – extent, duration, and depth. Using each source we compute the mean share of land under water in each hexagon for each month from July to December. This gives thirteen flood measures per hexagon: six monthly flood-shares from each satellite, and one depth. We take their first principal component.

3.5 Hydrological model inputs

To construct our instruments we use the LISFLOOD-FP hydrological model (Bates et al., 2010; Neal et al., 2012), which takes daily rainfall and topography (a digital map of ground elevation) as inputs. For topography we use FABDEM, which strips buildings and forest canopy from the Copernicus surface model, leaving the ground surface over which water flows, including drainage channels and irrigation canals. FABDEM does not capture large-scale flood defences, but these are largely absent in our setting. We take realised rainfall from NASA’s IMERG satellite product, since Pakistan has few ground weather stations, and long-range forecasts from the ECMWF seasonal ensemble.

3.6 Trade, transport cost and other model data

To estimate trade costs we use the arrivals recorded by Punjab’s market committees, which report daily to the provincial government the quantity of each commodity arriving in their market and the place it came from. This gives us bilateral trade flows between markets, which we use to estimate a gravity equation. To measure travel times we take the road network from OpenStreetMap, with road class from Geofabrik and surface quality from HOTOSM. We assign each class a speed, using published maximum speeds by road type and recorded average speeds. We then compute bilateral travel times between hexagon centres with Dijkstra’s algorithm. We take the population of each location from WorldPop 2020, which distributes census counts according to where satellites detect buildings, and the housing endowment from the Global Human Settlement Layer, which records the square metres of residential building in each 100m cell.

4 RESEARCH DESIGN

4.1 Endogeneity challenge

Our objective is to estimate the causal effect of flooding on household outcomes, using cross-sectional data collected one year after the floods.

$$Y_i = \beta F_i(k) + \alpha_d + \varepsilon_i \quad (1)$$

$$Y_i = \beta_{\text{own}} F_i^{\text{own}}(c) + \beta_{\text{near}} F_i^{\text{near}}(c) + \alpha_d + \varepsilon_i \quad (2)$$

where Y_i is an outcome for household i , $F_i(k)$ is the mean flood index over the hexagons within k km of its home, $F_i^{\text{own}}(c)$ is the mean over the hexagons within c km, $F_i^{\text{near}}(c)$ is the mean over the hexagons between c and 15 km, α_d is a district fixed effect and ε_i collects unobserved determinants of the outcome. β captures the total effect of flooding within k km of the household. β_{own} captures the effect of flooding within c km of the household, holding fixed flooding from c to 15 km, and β_{near} captures the effect of flooding from c to 15 km, holding fixed flooding within c km.

OLS estimation of equations 1 and 2 will yield biased estimates if flooding is correlated with unobserved determinants of outcomes. Such correlation may arise for physical reasons, since topography shapes both flood risk and agricultural productivity, or for behavioural reasons, if households sort and invest in response to anticipated flooding.

In both cases, the bias arises because realised flooding is correlated with expected flooding, which is in turn correlated with these unobserved determinants. We therefore employ a research design that exploits only the deviation of a location's flooding in 2022 from its expected flooding.

4.2 IV approach

We use the recentered instrumental variable approach of [Borusyak and Hull \(2023\)](#), which is designed for settings in which bias arises from non-random exposure to exogenous shocks.

In our setting, the shocks we exploit are rainfall, and a location's exposure to them depends on the physical characteristics that govern how water flows and accumulates across space. Flooding in

one location can therefore depend on rainfall in many other locations, since rain that falls far away but upstream of a location eventually flows to it and contributes to flooding there.

In stylised form, we write flooding in hexagon j on day t as

$$W_{jt} = w_{jt}(\mathbf{r}_{1:t}; \mathbf{X}), \quad (3)$$

where $\mathbf{r}_{1:t}$ is rainfall in every location up to day t and $\mathbf{X}$ includes the characteristics of all locations that are fixed before the monsoon begins, observed or not, such as topography, drainage infrastructure and pre-monsoon water levels. We summarise daily flooding over the T days of the monsoon with a single index,

$$Flood_j = s(W_{j1}, \dots, W_{jT}), \quad (4)$$

which therefore depends on rainfall in all locations on all days, $\mathbf{r} \equiv \mathbf{r}_{1:T}$.

Average rainfall varies across locations due to permanent features such as latitude, distance from the ocean and topography, and across years due to slow-moving factors such as ocean temperatures, both of which are known before the monsoon begins. Within a season, however, the timing and location of rainfall are very difficult to forecast, because the atmosphere is chaotic and small differences in initial conditions can produce very different rainfall patterns. Deviations from typical rainfall are therefore unlikely to be correlated with shocks to household outcomes realised after the onset of the monsoon.

Locations nonetheless differ systematically in their exposure to these shocks, that is, in the flooding they experience for a given realisation of $\mathbf{r}$. Random variation in rainfall therefore does not generate random variation in flooding.

[Borusyak and Hull \(2023\)](#) show that the bias from non-random exposure to exogenous shocks can be avoided by controlling for expected treatment. In our application, this corresponds to controlling for expected flooding, $\mathbb{E}[Flood_j]$, where the expectation is taken over all possible realisations of $\mathbf{r}$, given what is known before the monsoon begins.

We cannot compute this expectation directly, since we do not know the exact functional form of w_{jt} and in all likelihood observe only a subset of $\mathbf{X}$. Nor can we estimate it by averaging satellite-

observed flooding across past monsoons, since $\mathbf{X}$ may change over time, as embankments are built and channels silt up.

We instead use a hydrological model, which predicts flooding in each hexagon on each day as a function of rainfall and the subset of characteristics that we observe, $\hat{\mathbf{X}}$,

$$H_{jt} = H_{jt}(\mathbf{r}_{1:t}; \hat{\mathbf{X}}). \quad (5)$$

Following [Borusyak and Hull \(2025\)](#), we estimate the function $\hat{s}(\cdot)$ that combines these daily predictions into the best predictor of the satellite flood index, as described in Section 4.4,

$$\widetilde{Flood}_j(\mathbf{r}) = \hat{s}(H_{j1}, \dots, H_{jT}). \quad (6)$$

For each hexagon, we then compute the deviation of predicted flooding under 2022 rainfall from its expectation across possible monsoons,

$$Z_j = \widetilde{Flood}_j(\mathbf{r}^{2022}) - \mathbb{E}[\widetilde{Flood}_j(\mathbf{r})]. \quad (7)$$

We aggregate Z_j across hexagons to construct an instrument for each measure of flooding. The instrument for $F_i(k)$, $Z_i(k)$, is the mean of Z_j over the hexagons within k km of household i , and the instruments for $F_i^{\text{own}}(c)$ and $F_i^{\text{near}}(c)$, $Z_i^{\text{own}}(c)$ and $Z_i^{\text{near}}(c)$, are the means over the hexagons within c km and between c and 15 km.

Using these instruments, we estimate equations 1 and 2 by two-stage least squares. For equation 2, the two first-stage equations are

$$F_i^{\text{own}}(c) = \pi_1 Z_i^{\text{own}}(c) + \pi_2 Z_i^{\text{near}}(c) + \gamma_d + u_i, \quad (8)$$

$$F_i^{\text{near}}(c) = \pi_3 Z_i^{\text{own}}(c) + \pi_4 Z_i^{\text{near}}(c) + \delta_d + v_i, \quad (9)$$

and the second stage is

$$Y_i = \beta_{\text{own}} \widehat{F}_i^{\text{own}}(c) + \beta_{\text{near}} \widehat{F}_i^{\text{near}}(c) + \alpha_d + \varepsilon_i, \quad (10)$$

where $\widehat{F}_i^{\text{own}}(c)$ and $\widehat{F}_i^{\text{near}}(c)$ are the fitted values from the first stage.

4.3 Identifying assumptions

Let Ω denote all information available before the monsoon begins, including $\mathbf{X}$ and slow-moving factors, such as ocean temperatures, that shift the distribution of rainfall. 2SLS yields consistent estimates under two assumptions.

Assumption 1 (Correct rainfall distribution). *The distribution over which we take the expectation in equation 7 is the true distribution of rainfall given Ω ,*

$$\mathcal{R}(\cdot \mid \Omega) = \mathcal{R}_0(\cdot \mid \Omega).$$

Assumption 2 (Conditionally exogenous rainfall). *Which monsoon occurred, out of those that were possible given Ω , is unrelated to unobserved determinants of household outcomes,*

$$\mathbb{E}[\varepsilon_i \mid \mathbf{r}, \Omega] = \mathbb{E}[\varepsilon_i \mid \Omega].$$

Assumption 2 does not require rainfall to be unconditionally randomly assigned across locations. To see this, decompose the error into a component determined before the monsoon and a component realised after its onset, $\varepsilon_i = \mathbb{E}[\varepsilon_i \mid \Omega] + v_i$. The assumption is then equivalent to $\mathbb{E}[v_i \mid \mathbf{r}, \Omega] = 0$. The predetermined component may be correlated with topography, infrastructure or expected flooding. What is required is only that deviations of realised rainfall from what was expected given Ω are unrelated to shocks to household outcomes realised after the onset of the monsoon.

Neither assumption requires the hydrological model to be correct. If its functional form is wrong or $\hat{\mathbf{X}}$ omits relevant characteristics, the instrument is still recentered over the true distribution of rainfall, so errors in the model weaken the first stage but do not bias the estimates. This would not hold if we instead controlled for the model's expected flooding, since the part of expected flooding that the model misses would remain in the error.

4.4 Implementation

We simulate flooding with LISFLOOD-FP (Bates et al., 2010; Neal et al., 2012), a hydrodynamic model designed for large areas. The model takes daily rainfall and topography as inputs (Section 3) and outputs daily flood depth and extent in each hexagon. We input rainfall for July and August and run the model for five months, so that it captures the period over which flood water drains. This yields 151 days of depth and extent, or 302 outputs per hexagon. We simulate flooding under 2022 rainfall and under the rainfall of each monsoon from 1998 to 2021. In doing so, we assume that these historical monsoons represent the distribution of monsoons that could have occurred in 2022. As an alternative, we use the 51 members of the ECMWF seasonal forecast ensemble initialised in May 2022, which condition on the state of the oceans and atmosphere before the monsoon.

To aggregate the hydrological outputs into a prediction of the satellite flood index, we could apply the formula we use to construct the index itself, the first principal component of the monthly averages (Section 3). The model’s dynamics may differ from those of real floods, however, so we instead treat the 302 outputs as candidate instruments and follow the optimal instrument approach of Borusyak and Hull (2025), which builds on Chamberlain (1987). We estimate $\hat{s}(\cdot)$ as the best predictor of the satellite index given the 302 outputs, using a random forest. Intuitively, the random forest selects the weight on each day, and on depth relative to extent, that best predicts the satellite index. To avoid over-fitting, we cross-fit (Chernozhukov et al., 2018). We assign hexagons to five folds, clustered by H3 resolution-6 hexagon of approximately 36 km² to account for spatially correlated errors, and predict each fold with a forest trained on the other four.

Since $\hat{s}(\cdot)$ is a fixed function of the model’s outputs, we apply it to the outputs from each historical monsoon to compute $\widetilde{Flood}_j(\mathbf{r})$ and hence the expectation in equation 7.

4.5 Inference

We report Conley (1999) standard errors, which allow errors to be spatially correlated up to a distance cutoff. Since our instrument is generated by a single rainfall realisation propagating through the hydrological network, we set the cutoff at the scale over which the instrument is spatially correlated. We estimate this scale from the instrument’s correlogram, which plots the correlation between each




Figure 1: First stage, by treatment radius and by own cutoff

Notes: Real expenditure regressions, all households, district fixed effects. A: F of the instrument for flooding within the radius. B: Kleibergen–Paap rk Wald F for own and nearby flooding on their two instruments. C: correlation of the own and nearby instruments, after residualising on district fixed effects. Conley variance throughout. Dashed line at 10. $N = 4,536$ households.

hexagon’s instrument and those of its neighbours at increasing distances, and set the cutoff where the correlation falls to zero. The cutoff varies with the sample and the measure of flooding, and we report it in each table.

Since all households share one realisation of rainfall, we also conduct randomisation inference. We construct placebo instruments by treating each historical monsoon as if it were the realised one, and compare our estimates with the distribution of estimates under these placebos.

4.6 First stage results

Figure 1 summarises the first stage for each specification we estimate. Panel A plots the first-stage F -statistic for equation 1 at each radius k . Panel B plots the Kleibergen–Paap statistic for equation 2 at each cutoff c . The instruments are strong across all specifications, with statistics ranging from 10 to 147. Panel C plots the correlation between $Z_i^{\text{own}}(c)$ and $Z_i^{\text{near}}(c)$ within districts at each cutoff c , which determines how much independent variation we have to identify β_{own} and β_{near} .

5 EMPIRICAL RESULTS

5.1 Flooding in the surrounding area reduces consumption, conditional on flooding at home

We begin with our key result. Flooding reduces household consumption one year after the floods, and the effect grows with the size of the area around the household over which we measure flooding. Decomposing the effect, most of it is due to flooding beyond the household’s immediate surroundings rather than flooding close to its home.

We measure consumption as real expenditure, that is, household expenditure in the previous month, adjusted for household composition and deflated by a price index constructed from the food prices that households report paying (Section 3). Panel A of Figure 2 reports estimates of equation 1 for radii k from 0.5 to 15km. We scale each coefficient by the difference between the 10th and 90th percentiles of the within-district distribution of $F_i(k)$, so that it gives the effect of moving a household from a lightly to a heavily flooded area, at that radius.

The effect increases with the radius. Flooding within 0.5km reduces real expenditure by 4.9%, although this is not statistically significant ($p = 0.22$). Flooding within 1km reduces it by 6.8% ($p = 0.05$), rising to 9.7% within 3km and 11.2% within 5km. Beyond 5km the effect is stable, at 9.4% within 10km ($p = 0.02$) and 10.1% within 15km ($p = 0.01$).

Panel B decomposes the 15km effect into the contributions of flooding close to the home and flooding further away. Moving a household from the 10th to the 90th percentile of flooding within 15km raises both its own flooding, within c km, and its nearby flooding, between c and 15km. We compute the average change in each component that accompanies this move, $\Delta F^{\text{own}}(c)$ and $\Delta F^{\text{near}}(c)$,⁵ and multiply each by its coefficient from equation 2. For each cutoff c , Panel B plots the resulting contributions,

$$\underbrace{\beta_{\text{own}} \Delta F^{\text{own}}(c)}_{\text{own flooding}} \quad \text{and} \quad \underbrace{\beta_{\text{near}} \Delta F^{\text{near}}(c)}_{\text{nearby flooding}}. \quad (11)$$

⁵We regress own and nearby flooding separately on flooding within 15km, with district fixed effects. Each coefficient gives the average change in that component per unit of flooding within 15km, which we multiply by the 10th to 90th percentile gap in flooding within 15km.




Figure 2: Effect of flooding on real expenditure, by treatment radius and by own cutoff

Notes: IV estimates, district fixed effects. Outcome: log real expenditure per capita, all households. A: flooding within the radius, per within-district p10 to p90 of the regressor, times 100. B: the 15 km effect split into the part through own flooding within the cutoff (blue) and through nearby flooding from the cutoff to 15 km (orange), each coefficient times the move in that regressor along the 15 km p10 to p90, so the two add up to the dashed line, the 15 km total from A. Per-unit coefficients in Table A.2. 90 percent Conley intervals, Bartlett kernel. $N = 4,536$ households.

The two contributions sum approximately to the total effect within 15km from Panel A, shown by the dashed line.

Most of the effect is due to nearby flooding. Flooding more than 1km from the home accounts for 89% of the fall in real expenditure, flooding more than 5km away for 64%, and flooding more than 10km away for 49%. The contribution of own flooding is small and statistically insignificant at narrow cutoffs, and rises only as the cutoff widens to include flooding several kilometres away.

5.2 Why does flooding several kilometres from home reduce consumption?

And why does flooding close to the home have relatively modest impacts? We consider three explanations. We first show that the pattern is not driven by error in measuring household locations. The allocation of aid explains part of the small effect of own flooding, but not the large effect of nearby flooding. Instead, the evidence points to spillovers through labour, product and credit markets.

5.3 Candidate explanation 1: measurement error

Two types of measurement error could affect our estimates. The first is error in measured flooding at a location, $Flood_j$, and the second is error in the recorded location of each household.

Our IV approach already addresses the first. The flood index combines several satellite measures, which reduces noise, and the instrument is an independent prediction of flooding, so classical error in the index should not bias our estimates.

The instrument does not address the second. Suppose we record household i at $g_i = \ell_i + \xi_i$, where ℓ_i is its true location. Its measured own flooding, $F_i^{\text{own}}(c)$, then contains an error approximately equal to the displacement ξ_i times the spatial gradient of flooding at ℓ_i , which is largest at the edges of flooded areas. Since we construct the instrument at the same recorded location, this error is correlated with the instrument. Nearby flooding then partly proxies for true own flooding, so location error biases β_{own} downwards and β_{near} upwards (Huber, 2023).

We test for this in two ways. First, our results are robust to each of several measures of household location. Second, we use physical damage as a placebo outcome, since flooding several kilometres away cannot damage a household’s home or crops. If location error inflated the effect of nearby flooding, nearby flooding would appear to cause damage.

Figure 3 shows that it does not. With a 1km cutoff, own flooding raises the share of buildings damaged by 5.9 percentage points and lowers crop greenness by 33%, whereas nearby flooding has small and insignificant effects on both. Since location error biases own and nearby effects together, this also suggests that we do not substantially underestimate the effect of own flooding. The consumption effects of nearby flooding are therefore unlikely to be an artefact of location error.

Our estimand is nonetheless broader than the effect of flooding at a household’s home. Moving a household from the 10th to the 90th percentile of flooding within 1km raises the share of buildings damaged by 5.9 percentage points. Scaling our own-flooding estimates by this change would give an illustrative effect of home damage on consumption, although flooding also affects consumption through channels other than housing damage.




Figure 3: Effect of flooding on housing damage and crop greenness, by treatment radius and by own cutoff

Notes: IV estimates, district fixed effects. Top: share of buildings damaged, hexagons with at least 50 buildings, controlling for log buildings, building-weighted, $N = 2,497$ to $2,498$ hexagons. Bottom: September 2022 NDVI, cropland hexagons, controlling for the 2016 to 2020 September mean, coefficients as per cent of that mean's sample average (0.53), $N = 4,072$ to $4,075$ hexagons. A: flooding within the radius. B: own flooding within the cutoff (blue) and nearby flooding from the cutoff to 15 km (orange), from one regression. Both per within-district p10 to p90 of the regressor, times 100. 90 percent Conley intervals, Bartlett kernel.

5.4 Candidate explanation 2: aid

A second concern is that our estimates are net of aid. If relief targeted physically exposed households, it could offset the effect of own flooding, and so explain why flooding close to the home has little effect on consumption. We do not have random variation in aid. We can, however, estimate the effect of flooding separately for the two-thirds of households who report receiving no relief.




Figure 4: Effect of flooding on real expenditure, households reporting no relief

Notes: IV estimates, district fixed effects, households reporting no relief: the flooding coefficient from one regression per point on the pooled sample with flooding, a relief indicator and their interaction, instruments the recentred prediction and its interaction. Outcome: log real expenditure per capita. A: flooding within the radius, per within-district p10 to p90 of the regressor, times 100. B: the 15 km effect split into the part through own flooding within the cutoff (blue) and through nearby flooding from the cutoff to 15 km (orange), each coefficient times the move in that regressor along the 15 km p10 to p90, so the two add up to the dashed line, the 15 km total from A. Per-unit coefficients in Table A.2. C: the own part of B over the 15 km total, for all households (pink) and households reporting no relief (pale pink). 90 percent Conley intervals, Bartlett kernel. $N = 4,178$ households.

Figure 4 repeats Figure 2 for households who report receiving no relief. For these households, the effect of flooding is larger at every radius, reaching 15% within 15km. Own flooding also accounts for a larger share of the effect. Panel C compares the share of the 15km effect due to own flooding for all households and for households without relief. With a 1km cutoff, this share rises from 11% to 24%, and it is higher for households without relief at every cutoff. For households who received relief, the effect of flooding is smaller and is driven almost entirely by nearby (not own) flooding (Appendix Figure 12).

Together, these results suggest that households with more own flooding received more relief, which helped them maintain consumption, whereas relief did not flow to households on the basis of flooding in the surrounding area. Our estimates for the full sample are therefore likely to understate the effect of own flooding, and so to overstate the share of the effect due to nearby flooding. Aid does not, however, explain away the effect of nearby flooding. Even among households who received no

relief, flooding more than 1km away accounts for 76% of the effect, and flooding more than 5km away for 58%.

5.5 Candidate explanation 3: spillovers

Even after accounting for measurement error and aid, the effects of own flooding on consumption remain small relative to those of flooding in the surrounding area. We argue next that this pattern reflects spillovers through markets.

Suppose that the area within 1km of each household were a self-contained village existing in autarky. Its residents produce all the food and goods they consume, and do not trade, commute, borrow or lend outside the village. Only flooding within 1km would then affect their consumption, and β_{near} would be zero.

If villages are instead connected through goods, labour and credit markets, a flood in one village can affect people living in others. A household may, for example, lose its job in a flooded village, be unable to reach its usual market, or need to support flooded relatives. More generally, a flood in one location can reduce the demand for labour or the supply of food in a local market, and so change wages and prices in many other locations. Integration therefore shrinks β_{own} , since the costs of flooding in a household's own village are shared with people elsewhere, providing partial insurance against local shocks. Conversely, β_{near} becomes non-zero, since the household in turn shares in the costs of flooding in other villages.

In the limit in which everything that a household consumes is traded and there is no cost of commuting or transporting goods both β_{own} and β_{near} would be zero, since only flooding across the economy as a whole would matter, and its effects would be absorbed by our district fixed effects. When trade and commuting are costly, and these costs rise with distance, spillovers are concentrated in space, so flooding in the area around a household's home becomes an important determinant of its outcomes. In the model of Section 6, consumption loss is determined by the flood-induced change in market access, which is effectively a distance-weighted average of flooding in the surrounding area.

If spillovers operate through markets, they should affect the quantities and prices determined in those markets. Spillovers through labour markets should affect employment, wages and commuting, those through goods markets should affect prices, and those through credit markets should affect borrowing and asset sales. We consider each in turn, starting with the labour market.

Figure 5 shows the effect of flooding on wage employment, measured as the share of working-age men in the household who work for others rather than for themselves. Panel A shows that flooding reduces wage employment, and that the effect grows with the radius over which we measure flooding. Flooding within 1km reduces wage employment by 1.5 percentage points ($p = 0.60$), rising to 3.4ppts within 5km and 8.9ppts within 15km ($p = 0.01$). The effect is driven by households whose men worked in agriculture before the floods, for whom flooding within 15km reduces wage employment by 16.9 percentage points, compared with 0.9ppts for other households (Appendix B.1).

Panel B shows that the effect is due almost entirely to nearby flooding. With a 1km cutoff, nearby flooding accounts for a fall of 9.3 percentage points, whereas own flooding slightly increases wage employment, by 0.5 percentage points ($p = 0.26$). The pattern persists as the cutoff widens. With a 5km cutoff, nearby flooding accounts for a fall of 11.0 percentage points and own flooding for a rise of 2.2 ($p = 0.16$). This pattern is consistent with households supplying more labour in response to their own losses, while flooding in the surrounding area reduces the demand for their labour.

Consistent with these effects operating through access to work, own flooding increases the probability that men work more than 1km from home, whereas nearby flooding reduces it (Appendix Figure 13). With a 1km cutoff, own flooding raises this probability by 2.2 percentage points ($p < 0.01$) and nearby flooding lowers it by 2.7 ($p = 0.21$). With a 5km cutoff, own flooding raises it by 4.8 percentage points ($p < 0.01$) and nearby flooding lowers it by 5.5 ($p = 0.01$). This suggests that households flooded at home but surrounded by dry land can seek work elsewhere, but that this becomes increasingly difficult as more of the surrounding area is flooded. These results are in line with Brooks and Donovan (2020), who find that bridges, by maintaining access to outside labour markets, protect earnings from local flooding.

Part of the fall in wage employment is offset by a rise in self-employment (Appendix Figure 14). Moving a household from the 10th to the 90th percentile of flooding within 15km raises the share




Figure 5: Effect of flooding on working for others and on the food price index

Notes: IV estimates, district fixed effects, all households. Top: share of the household's men working for others, $N = 4,478$ households. Bottom: log household food price index, $N = 4,536$ households. A: flooding within the radius, per within-district p10 to p90 of the regressor, times 100. B: the 15 km effect split into the part through own flooding within the cutoff (blue) and through nearby flooding from the cutoff to 15 km (orange), each coefficient times the move in that regressor along the 15 km p10 to p90, so the two add up to the dashed line, the 15 km total from A. Per-unit coefficients in Table A.2. 90 percent Conley intervals, Bartlett kernel.

of its men in self-employment by 4.9 percentage points ($p = 0.01$), so that total employment falls by only 3.2 percentage points ($p = 0.10$). As with wage employment, the rise in self-employment is driven by nearby flooding, which accounts for 5.0 percentage points with a 1km cutoff. This is consistent with labour rationing, whereby workers displaced from wage work turn to lower-return forms of self-employment (Breza et al., 2021). We find no effect of flooding on daily wages. This is in line with wages being rigid downwards in village labour markets (Kaur, 2019), so that a fall in the demand for labour reduces employment rather than wages.

Figure 5 shows the effect of flooding on food prices, measured using a household price index constructed from the prices that households report paying for food (Section 3). Panel A shows that flooding raises food prices, and that the effect grows with the radius over which we measure flooding. Flooding within 1km raises food prices by 4.4% ($p = 0.07$), rising to 6.2% within 5km and 8.1% within 15km ($p = 0.02$). Panel B shows that this effect is due almost entirely to nearby flooding. With a 1km cutoff, nearby flooding raises food prices by 7.6%, whereas own flooding has no effect (0.6%, $p = 0.25$). Even at 10km, flooding beyond the cutoff accounts for the entire effect. This is consistent with food prices being set in markets that serve a wider area than a single village, so that what matters for the price a household pays is the loss of supply across that area, rather than in its own village. This is in line with [Allen and Atkin \(2022\)](#), who show that highway expansion in India made local prices less responsive to local yields and more responsive to yields elsewhere.

Figure 6 shows the effect of flooding on borrowing and asset sales, measured by whether the household has taken a loan and whether it has sold household assets since the floods. Panel A shows that flooding reduces borrowing. Flooding within 1km reduces the share of households with a loan by 3.0 percentage points ($p = 0.05$), rising to 5.8 within 5km and 4.9 within 15km ($p = 0.03$). Panel B shows that, at narrow cutoffs, this effect is due almost entirely to nearby flooding. With a 1km cutoff, nearby flooding reduces borrowing by 4.5 percentage points, whereas own flooding has no effect (0.5, $p = 0.32$). Flooding also raises asset sales, by 2.7 percentage points within 15km ($p = 0.05$), again almost entirely because of nearby flooding, which accounts for 2.4 percentage points with a 1km cutoff and 2.4 with a 10km cutoff. This suggests that households flooded at home can continue to borrow from lenders in the surrounding area, but that when the surrounding area is also flooded, credit becomes scarce and households instead sell assets. The rise in asset sales offsets only part of the fall in borrowing, possibly because assets are also harder to sell when many nearby households are affected.




Figure 6: Effect of flooding on loans and asset sales, by treatment radius and by own cutoff

Notes: IV estimates, district fixed effects, all households. Top: household has taken a loan, $N = 4,683$ households. Bottom: household sold assets, $N = 4,683$ households. A: flooding within the radius, per within-district p10 to p90 of the regressor, times 100. B: the 15 km effect split into the part through own flooding within the cutoff (blue) and through nearby flooding from the cutoff to 15 km (orange), each coefficient times the move in that regressor along the 15 km p10 to p90, so the two add up to the dashed line, the 15 km total from A. Per-unit coefficients in Table A.2. 90 percent Conley intervals, Bartlett kernel.

6 A SPATIAL MODEL OF DISASTER INCIDENCE

Our reduced-form results show that a household's losses depend on flooding both close to its home and in the surrounding area, so the incidence of a flood, meaning the distribution of welfare losses across locations, need not coincide with its physical footprint.

The reduced form cannot, however, recover the full, general-equilibrium incidence of a flood. Our estimates capture only flooding within 15km, whereas effects may propagate further, through

trading partners or general-equilibrium adjustments in wages and prices. Moreover, district fixed effects absorb any effects common to all households in a district, including district- or province-wide general-equilibrium effects (Nakamura and Steinsson, 2014; Chodorow-Reich, 2020; Adão et al., 2019; Wolf, 2023).

We therefore build a spatial model in which many locations are linked through costly trade, following Eaton and Kortum (2002), and costly commuting, following the quantitative spatial literature (Ahlfeldt et al., 2015). The costs of a flood are shared through goods and labour markets, the two main channels through which we find spillovers in the data. Taking a stand on how shocks propagate allows us to recover the welfare loss in every location and to evaluate counterfactual floods and aid policies.

6.1 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 and sector 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 and Kortum (2002), with a few departures. There are two sectors, agriculture A and non-agriculture M , each with a continuum of goods which can be produced anywhere. Agricultural goods are produced from labour and land using Cobb–Douglas technology with labour share α ; non-agricultural goods are produced from labour alone. Regions differ in their average productivity and their idiosyncratic productivity for each good. Productivity for good u in sector s in location j is $z_j^s(u)$ – drawn from a Fréchet distribution with scale parameter T_j^s and shape parameter θ_T . T_j^s captures the absolute advantage of a location in that sector, whilst θ_T governs dispersion across goods. A location can be a strong producer of food and a weak producer of everything else. The unit cost of producing good u in sector s in location j is:

$$\frac{c_j^s}{z_j^s(u)} \quad \text{with} \quad c_j^A \propto (w_j^A)^\alpha r_j^{1-\alpha}, \quad c_j^M = w_j^M, \quad (12)$$

where w_j^s is the sector- s 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 split their spending between the two sectors in Cobb–Douglas shares, with agricultural share χ , and have identical CES preferences over the continuum of goods within each sector, with elasticity of substitution σ . They can buy goods from anywhere, but shipping is subject to iceberg trade costs $\kappa_{ij}^G = t_{ij}^G$, rising in the travel time t_{ij} between the two locations, 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 i 's sector- s spending that goes to j is:

$$\pi_{ij}^{G,s} = \frac{T_j^s (c_j^s \kappa_{ij}^G)^{-\theta_T}}{\sum_k T_k^s (c_k^s \kappa_{ik}^G)^{-\theta_T}}, \quad (13)$$

All else equal, i will buy more from j if it has higher productivity, lower wages and land rents, or is closer. Trade costs and the elasticity are common to both sectors, but the two trade networks differ because productivities and bundle costs do. The sector price index in i falls with its goods-market access in that sector: $P_i^s \propto (\text{GMA}_i^s)^{-1}$, where $\text{GMA}_i^s \equiv [\sum_k T_k^s (c_k^s \kappa_{ik}^G)^{-\theta_T}]^{1/\theta_T}$. The cost of living combines the two, $P_i = (P_i^A)^\chi (P_i^M)^{1-\chi}$.

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

$$\pi_{ij}^s = \frac{(w_j^s / \tau_{ij})^{\theta_L}}{\sum_k \sum_t (w_k^t / \tau_{ik})^{\theta_L}}, \quad (14)$$

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. Both sectors at j face the same commuting cost, so distance factors out of the sector choice: summing over sectors gives the share who commute to j ,

$$\pi_{ij} = \frac{(\mathcal{W}_j / \tau_{ij})^{\theta_L}}{\sum_k (\mathcal{W}_k / \tau_{ik})^{\theta_L}}, \quad \mathcal{W}_j \equiv \left[\sum_s (w_j^s)^{\theta_L} \right]^{1/\theta_L}, \quad (15)$$

where $\mathcal{W}_j$ is the wage index at j , and dividing gives the sector split $\pi_{ij}^s/\pi_{ij} = (w_j^s/\mathcal{W}_j)^{\theta_L}$. The sector composition of the workforce at j is therefore the same for every origin, and the sectoral wage gap at j is pinned down by the sectoral employment split, $w_j^M/w_j^A = (L_j^M/L_j^A)^{1/\theta_L}$. The expected wage is $\bar{w}_i = \sum_j \sum_s \pi_{ij}^s w_j^s$. 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 $\text{CMA}_i \equiv \left[\sum_j (\mathcal{W}_j/\tau_{ij})^{\theta_L} \right]^{1/\theta_L}$, so that $\bar{v}_i \propto \text{CMA}_i$.

Factor markets and incomes. Firm revenue in sector s in location j is total spending on its goods, $X_j^A = \chi \sum_i \pi_{ij}^{G,A} E_i$ and $X_j^M = (1 - \chi) \sum_i \pi_{ij}^{G,M} E_i$, where E_i is income in i and $\pi_{ij}^{G,s}$ its sector- s expenditure share on j . We assume factor markets are perfectly competitive, so revenue is split in fixed Cobb–Douglas shares: agricultural workers earn $w_j^A L_j^A = \alpha X_j^A$ and the landowner earns $r_j H_j = (1 - \alpha) X_j^A$, whilst non-agricultural workers take all of X_j^M , so $w_j^M L_j^M = X_j^M$, where $L_j^s = \sum_i L_i \pi_{ij}^s$ is the workforce commuting in to that sector. Land rents come out of agricultural revenue alone, since non-agriculture uses no land.

Welfare. A worker spends labour income on goods and enjoys their 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 job preference draws) is therefore

$$V_i^w = \left(\frac{\bar{v}_i}{P_i} \right)^{1-\beta} B_i^\beta, \quad (16)$$

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 in both sectors (CMA_i) and to cheap food and cheap non-food (GMA_i^A and GMA_i^M). The landowner has no housing and consumes only goods, so their expected utility is $V_i^r = (1 - \alpha) X_i^A / P_i$.

6.2 Equilibrium

Given location fundamentals $\{T_i^A, T_i^M, H_i, B_i, L_i\}_{i=1}^N$, parameters $\{\alpha, \beta, \chi, \theta_T, \theta_L, \sigma\}$, and bilateral trade and commuting costs $\{\kappa_{ij}^G, \tau_{ij}\}$, a competitive equilibrium is a vector of factor prices $\{w_j^A, w_j^M, r_j\}_{j=1}^N$

at which agents solve the problems above – pinning down c_j^s , $\pi_{ij}^{G,s}$, P_i^s , P_i , π_{ij}^s , $\bar{w}_i$ and $\bar{v}_i$ – and all markets clear:

- (i) Labour: each sector's wage bill equals labour's share of that sector's revenue, $w_j^A L_j^A = \alpha X_j^A$ and $w_j^M L_j^M = X_j^M$, where $L_j^s = \sum_i L_i \pi_{ij}^s$ is the workforce commuting in to that sector.
- (ii) Land: rents are pinned by land's share of agricultural revenue, $r_j = (1 - \alpha) X_j^A / H_j$.
- (iii) Goods: each sector's revenue equals spending on its output, $X_j^A = \chi \sum_i \pi_{ij}^{G,A} E_i$ and $X_j^M = (1 - \chi) \sum_i \pi_{ij}^{G,M} E_i$, where income is $E_i = L_i \bar{w}_i + (1 - \alpha) X_i^A$.

We fix the agricultural wage in the first location as the numeraire, $w_1^A = 1$. By Walras's law one market-clearing condition is then redundant.

6.3 A flood in the model

We define a flood as a simultaneous negative shock to agricultural productivity and housing. Locations which are flooded experience a reduction in T_j^A , lowering average productivity in agriculture, 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 agricultural productivity, on the other hand, can have impacts that ripple out across the entire economy. They occur at every production location at once, and 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.

$$d \ln V_i^w = \underbrace{\beta d \ln B_i}_{\text{housing}} + (1 - \beta) \left[\underbrace{\sum_j \sum_s \pi_{ij}^s d \ln w_j^s}_{\text{wage}} - \underbrace{\sum_j \sum_s \chi_s \pi_{ij}^{G,s} d \ln p_j^s}_{\text{cost of living}} \right], \quad (17)$$

where $\chi_A = \chi$ and $\chi_M = 1 - \chi$. 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 and sectors $d \ln w_j^s$ weighted by the probability that they held each job prior to the flood π_{ij}^s , and the change in prices in all locations and sectors $d \ln p_j^s$ weighted by the share of their expenditure that was spent on goods from j in sector s prior to the flood.

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

$$d \ln V_i^r = \underbrace{d \ln r_i}_{\text{rent}} - \underbrace{\sum_j \sum_s \chi_s \pi_{ij}^{G,s} d \ln p_j^s}_{\text{cost of living}}, \quad d \ln r_i = d \ln X_i^A, \quad (18)$$

Together these equations map any flood to its welfare cost at every location.

The wage and price changes are themselves general-equilibrium responses to the flood. To see where they come from, recall that the price in a location is pinned down by unit cost. Let p_j^A be the mill price of an average agricultural good produced in j . Firms are perfectly competitive, so they price at marginal cost – the input bundle c_j^A per unit of output:

$$p_j^A \propto \frac{c_j^A}{z_j^A}, \quad c_j^A = (w_j^A)^\alpha r_j^{1-\alpha}, \quad (19)$$

A flood lowers T_j^A , hence z_j^A : 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,

$$\underbrace{-a_j}_{\text{unit cost up}} = \underbrace{d \ln p_j^A}_{\text{buyers pay more}} - \underbrace{\alpha d \ln w_j^A}_{\text{workers earn less}} - \underbrace{(1 - \alpha) d \ln r_j}_{\text{landowners earn less}}, \quad (20)$$

where a_j is the proportional change in average agricultural 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.

The trade and commuting elasticities govern how the loss is divided between them. 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 mobile, the agricultural wage is held up by workers leaving for jobs elsewhere or in non-agriculture, and the residual loads onto land rent. The loss falls on whichever margin cannot adjust.

6.4 Parameter estimation

The trade elasticity and trade costs: We estimate the trade elasticity and elasticity of trade costs to travel time in two steps. To estimate the product $\theta_T \tau^G$ we run a gravity regression. We parameterise iceberg trade costs as a constant-elasticity function of travel time, $\kappa_{ij}^G = t_{ij}^{\tau^G} e_{ij}$, and defining origin and destination fixed effects λ_i and γ_j we can write the trade share (13) as:

$$\pi_{ij}^{G,s} = \exp(\lambda_i + \gamma_j) t_{ij}^{-\theta_T \tau^G} \eta_{ij}, \quad (21)$$

We estimate (21) using internal trade data covering arrivals at all the major market places in Punjab (described in Section 3.6), and travel times by road as described in Section 3. To address zero flows, we estimate by Poisson pseudo-maximum likelihood. Using pooled data for the trade years 2018/19 to 2024/25, we get $\theta_T \tau^G = 1.99$ with a standard error of (0.081).

To separate θ_T from τ^G , we use the fact that in the model the change in log agricultural prices in a location equals $-1/\theta_T$ times the change in its log agricultural market access. We regress the log food price index from our survey on the flood-induced change in agricultural market access implied by the 2022 floods. To address endogeneity, we use an IV approach analogous to that of Section 4. This gives $\theta_T = 2.70$, which implies $\tau^G = 1.99/\theta_T = 0.74$.

The commuting elasticity and commuting costs: We estimate the product $\theta_L \tau^L$ from a gravity regression of commuting flows between hexagons in our survey data on travel time, by Poisson

pseudo-maximum likelihood, with a dummy for working in the home hexagon. This gives $\theta_L \tau^L = 2.36$ (0.08). To separate the two, we choose θ_L to match the standard deviation of log wages across workplaces in our survey, 0.59, which gives $\theta_L = 2.79$ and $\tau^L = 0.85$. We calibrate the cost of working outside the home hexagon to match the 77% of workers who work in their own hexagon.

The labour share. We set $\alpha = 0.40$, the share of agricultural revenue paid to workers rather than to landowners. We estimate it from the Agriculture Policy Institute’s 2021–22 cost-of-production data, which survey growers and traders. This is close to the value of 0.46 that [Sotelo \(2020\)](#) uses for Peru.

The housing share. We set $\beta = 0.19$, the share of housing in utility, which governs how much of a flood’s welfare cost arises through destroyed housing rather than through prices and wages. We compute it as the housing share of spending among rural households in the 2018–19 Pakistan Social and Living Standards Measurement Survey. This is slightly below typical values in the literature, including the value of 0.25 that [Sturm et al. \(2021\)](#) obtain for Dhaka using the same approach.

The elasticity of substitution across goods. We take $\sigma = 2.4$ from [Sotelo \(2020\)](#). This parameter does not affect our results, since it governs substitution across goods within a sector, and a flood affects every agricultural good in a location equally. What matters instead is substitution across sources, which is governed by θ_T .

6.5 Quantifying the incidence of a flood

Given a map of flooding, the model delivers the change in the average welfare of residents in each location. We compute this in four steps. First, we define the extent of the economy and the unit of a location. Second, we invert the model for the unobserved pre-flood productivities. Third, we estimate the effects of flooding on agricultural productivity and housing. Fourth, we combine a flood map with these elasticities to compute post-flood productivities and housing, and solve for the new equilibrium.

Step 1: Geography. We quantify the model for the province of Sindh, which we divide into 14,037 hexagons at resolution 7 of the H3 grid, each with an area of approximately 5km² and a radius

of about 1.2km. Each hexagon is a location in the model. We treat Sindh as a closed economy, with no trade or commuting with the rest of Pakistan.

Step 2: Model inversion. We recover the unobserved pre-flood productivities by inverting the model, using data on agricultural revenue X_j^A , residential population L_i and the share of each hexagon’s residents working in agriculture (Section 3). Agricultural wages satisfy labour market clearing,

$$w_j^A L_j^A = \alpha X_j^A,$$

where the agricultural workforce L_j^A depends on wages everywhere through commuting. Given agricultural wages, the agricultural employment share pins down non-agricultural wages. If few residents work in agriculture despite well-paid agricultural jobs nearby, non-agricultural jobs within reach must pay well. Wages and land rents give each hexagon’s income E_i . Productivities then satisfy goods market clearing,

$$X_j^s = \chi_s \sum_i \pi_{ij}^{G,s} E_i,$$

so a hexagon that earns high revenue despite high unit costs or poor access to buyers must be productive.

Step 3: Effects of flooding. To simulate a flood in the model, we need to know how it changes each location’s fundamentals. By assumption, a flood changes only agricultural productivity and the housing stock, and we estimate the elasticity of each with respect to flood intensity using the 2022 floods. For productivity, we estimate a difference-in-differences regression of agricultural yields, output per km² of cropland, on flooding in each hexagon, using data from 2018 to 2022, and use the change in yields as a proxy for the change in productivity. For housing, we estimate the effect of flooding in each hexagon on the share of its buildings damaged, using the same instrumental-variable approach as in Section 4, and assume that a damaged house loses 95% of its housing services.

Step 4: New equilibrium. For a given flood map, we compute post-flood agricultural productivity and housing in each hexagon from its flood intensity and the estimated elasticities. We then solve for the new equilibrium, holding parameters, trade and commuting costs, and residential population fixed. The incidence of the flood is the resulting change in the average welfare of each hexagon’s residents.

7 THE ECONOMIC INCIDENCE OF SINDH'S 2022 FLOODS

We use the model to evaluate the economic incidence of the 2022 floods in Sindh. Panel A of Figure 7 is a map of the 2022 floods in Sindh, measured using the same flood index as in our reduced-form analysis. Two in five hexagons experienced almost no flooding. For the median hexagon, 1.5% of land was under water in September and 0.2% in December, with a maximum depth of 0.15m. For hexagons at the 75th, 90th and 95th percentiles of the flood index, the corresponding figures were 15%, 43% and 66% in September, 1.7%, 4.8% and 11.8% in December, and 0.58m, 1.14m and 1.11m.

Our estimates imply that these floods reduced average agricultural productivity across Sindh by 14.1% and average housing services by 10.9%. For the median hexagon, productivity fell by 4.7% and housing services by 3.3%. At the 75th, 90th and 95th percentiles, the corresponding falls were 19.1%, 41.9% and 62.2% for productivity and 13.9%, 31.8% and 49.7% for housing services.

Panel B of Figure 7 shows the resulting change in welfare in each hexagon. Welfare fell everywhere, including in hexagons that experienced almost no flooding. The average loss was 12.2%, and




Figure 7: The 2022 flood and the model-predicted welfare loss

Notes: 13,416 hexagons of about 5 square kilometres in the twenty crop-classified districts; districts outside the model are unfilled. A: the flood index, the first principal component of monthly MODIS and Sentinel-1 inundation shares and FLEXTTH depth, shifted so the least flooded hexagon is zero. B: the model-predicted fall in welfare, housing included, per cent of the no-flood level. Colour scales are square-root.

the smallest was 1.6%. For the median hexagon, welfare fell by 8.3%, and at the 75th, 90th and 95th percentiles by 14.8%, 24.3% and 34.2%.

Using the model, we can decompose each hexagon's welfare loss into the part due to its own flooding and the part due to flooding elsewhere. Let f_j denote flooding in hexagon j and f_{-j} flooding in all other hexagons, so that welfare in j is $V_j(f_j, f_{-j})$. For each hexagon, we solve the model four times: without the flood, $V_j(0, 0)$; applying the 2022 flood map to all hexagons, $V_j(f_j, f_{-j})$; applying it to j only, $V_j(f_j, 0)$; and applying it to all hexagons except j , $V_j(0, f_{-j})$. We can then compute the total effect of the flood on j , $V_j(f_j, f_{-j}) - V_j(0, 0)$, the direct effect, $V_j(f_j, 0) - V_j(0, 0)$, and the indirect effect, $V_j(0, f_{-j}) - V_j(0, 0)$.

Panels A and B of Figure 8 show the direct and indirect effects for each hexagon. The direct effect follows the pattern of flooding and ranges from 0 to 59%. The indirect effect is smoother, higher in the north and lower in the south, and ranges from 0.8 to 25%. For the average hexagon, 73% of the welfare loss is due to flooding elsewhere. This share falls with own flooding, from 99% in the least flooded decile to 32% in the most flooded. Since housing losses are by assumption entirely driven by flooding where you live, the share of loss due to indirect exposure is higher for real consumption, at 80% on average, ranging from 99% in the least flooded decile to 46% in the most flooded.




Figure 8: Welfare loss from own flooding and from flooding elsewhere

Notes: 13,416 hexagons. A: the welfare loss each hexagon would have taken from its own flooding alone, housing included. B: the loss from flooding in every other hexagon, its own flooding set to zero. Both per cent of the no-flood level, on one square-root colour scale.

8 TARGETING DISASTER ASSISTANCE

We have shown that, for the 2022 floods, most of the welfare loss in the average location was due to flooding elsewhere. Disaster relief, by contrast, is typically targeted on physical exposure, such as flooding or damage to homes.

Do our results imply that relief is misallocated? Not necessarily. First, own flooding may be highly correlated with total loss, in which case targeting on it reaches the right locations anyway. Second, the 2022 floods were unusually large and dispersed, and own flooding may be a better guide to loss in smaller or more concentrated floods. Third, aid may have general-equilibrium effects. Aid that restores production in flooded locations would also reverse the losses those locations impose on others. We consider each in turn.

8.1 Is own flooding a good proxy for need?

Although own flooding accounts for only a minority of the average location's loss, it may still be a good guide to which locations lose most.

8.1.1 Simple targeting rules

A simple first way to assess this is to compare the two rankings. We rank hexagons by their own flooding, measured by the flood index, and by their model-predicted welfare loss. A regression of the loss rank on the flooding rank yields an R^2 of 0.71, so a hexagon's rank in flooding explains 71% of the variation in its rank in welfare loss.

This R^2 is difficult to interpret on its own. A more concrete test compares the allocations implied by simple targeting rules based on physical exposure and on model-predicted welfare loss. After a disaster, governments often declare a list of affected areas and pay a fixed transfer to households living in them. After the 2022 floods, for example, every household below the eligibility threshold of Pakistan's standing cash transfer programme (the Benazir Income Support Programme) that lived in an area notified as flood-affected was entitled to a one-off transfer of 25,000 rupees ([Ministry of Finance, Government of Pakistan, 2023](#); [Associated Press of Pakistan, 2022](#)).

Suppose the government could pay a transfer to residents of 40% of hexagons, and was choosing between paying the 40% most flooded and paying the 40% with the largest model-predicted welfare losses. Panel A of Figure 9 plots each hexagon's percentile in flooding against its percentile in welfare loss, with lines at the 40% cutoff. The two rules agree on 4,479 hexagons (top right). The flood rule would miss the remaining 888 hexagons with the largest model-predicted welfare losses, 17% of the total, because their flooding falls below the cutoff (top left). If the goal were to compensate welfare losses, these would be exclusion errors. The flood rule would instead pay 888 hexagons with flooding above the cutoff but welfare losses below it (bottom right), which would be inclusion errors.

Panel B repeats this comparison for every share of hexagons paid. It plots the share of the hexagons with the largest model-predicted welfare losses that the flood rule would miss. When the government pays 20% of hexagons, the flood rule misses 12% of the worst affected. When it pays half, the flood rule misses 18%. The share falls towards zero beyond this point, because both rules then pay almost every hexagon. Excluding housing losses, which are driven by own flooding by construction, the flood rule misses 18% and 22% of the worst affected. Own flooding is therefore a




Figure 9: Targeting on own flooding against targeting on model-predicted welfare loss

Notes: A: each point is one of 13,416 hexagons; horizontal axis the percentile of the 2022 flood index, vertical axis the percentile of the model-predicted welfare loss, housing included; lines mark the 40 percent cutoff on each axis. B: the share of the hexagons with the largest losses that a rule paying the same share of the most flooded hexagons does not pay, for the welfare loss and for the welfare loss excluding housing damage, computed at every 2.5 percent of hexagons paid.

better guide to loss among the most flooded hexagons than among the moderately flooded, and even a narrowly targeted flood rule misses around one in eight of the worst affected.

Figure 10 maps the two types of error. The exclusion errors are less-flooded hexagons in heavily flooded surroundings. The median hexagon in this group lies at the 47th percentile of own flooding but the 68th percentile of flooding in other hexagons within 10km. Flooding elsewhere accounts for 92% of its welfare loss. The inclusion errors are the opposite case, at the 65th and 44th percentiles, and flooding elsewhere accounts for 58% of their welfare loss.




Figure 10: Inclusion and exclusion errors under the flood rule

Notes: The four groups of Figure 9 on the map, at the 40 percent cutoff: hexagons both rules pay, hexagons with welfare loss above the cutoff but flooding below it (exclusion errors), hexagons with flooding above the cutoff but welfare loss below it (inclusion errors), and the rest.

8.1.2 Optimal targeting

To measure the welfare cost of inclusion and exclusion errors, we must make an assumption about social welfare. We set out the planner's problem, and use it to quantify the value of providing the planner with our model-based estimates of welfare loss.

We assume that the planner cares about the welfare losses caused by the flood, and places more weight on larger losses. The planner can make transfers to households in different locations from a fixed budget, and allocates them to minimise the sum, across households, of an increasing and convex function of each household's welfare loss. Let $V_{ih}(\mathbf{f}, \mathbf{y})$ denote the welfare of household h in hexagon i , given flooding $\mathbf{f}$ and transfers $\mathbf{y}$ in all hexagons. The planner solves

$$\min_{\{y_{ih} \geq 0\}} \sum_i \sum_h \omega_{ih} g(V_{ih}(\mathbf{0}, \mathbf{0}) - V_{ih}(\mathbf{f}, \mathbf{y})) \quad \text{s.t.} \quad \sum_i \sum_h y_{ih} = B, \quad (22)$$

where $\omega_{ih} > 0$ are welfare weights, g is strictly increasing and strictly convex, and the budget B is financed from outside the economy.

Two features of this problem deserve comment. First, the planner cares about losses rather than levels of welfare, reflecting the apparent aim of disaster relief to compensate households for

what they lost rather than to reduce poverty. For example, post-disaster needs assessments organise assistance around damage and loss (World Bank et al., 2013), and US federal disaster assistance has no income test (Webster, 2023). The two objectives imply the same allocation when pre-flood welfare is equal across households. Otherwise, the loss-based planner places more weight on richer households, who have more to lose. Second, the budget is financed from outside the economy.

If the planner observes each household's welfare loss, the optimal transfers satisfy

$$\sum_k \sum_{h'} \omega_{kh'} g'(V_{kh'}(\mathbf{0}, \mathbf{0}) - V_{kh'}(\mathbf{f}, \mathbf{y})) \frac{\partial V_{kh'}}{\partial y_{ih}} = \mu - \lambda_{ih}, \quad \lambda_{ih} \geq 0, \quad \lambda_{ih} y_{ih} = 0, \quad (23)$$

where μ is the multiplier on the budget constraint and λ_{ih} is the multiplier on the non-negativity constraint. The social value of a transfer to household h in hexagon i is the reduction in losses it causes in every hexagon, valued at g' .

If transfers are made in units of utility, so that a transfer of y_{ih} raises the welfare of its recipient by y_{ih} , and transfers have no general-equilibrium effects, then equation 23 simplifies to

$$\omega_{ih} g'(V_{ih}(\mathbf{0}, \mathbf{0}) - V_{ih}(\mathbf{f}, \mathbf{y})) = \mu - \lambda_{ih}. \quad (24)$$

With equal welfare weights, the planner gives nothing to households whose losses fall below a cutoff and reduces the losses of all other households to that cutoff, with the cutoff set to exhaust the budget.

We use this set-up to compare three allocations of the same budget. In each, the planner sets a per-household transfer for each hexagon. In the first, this transfer is proportional to the hexagon's own flooding. In the second, it depends only on own flooding, but the planner chooses the schedule optimally using the predicted welfare losses. In the third, the planner sets it using the hexagon's predicted welfare loss, according to equation 24. Comparing the first and second allocations isolates the value of information about welfare losses. Comparing the second and third isolates the value of conditioning transfers on that information.

Take, for example, a planner with $g(\ell) = \ell^{1+\varepsilon}/(1+\varepsilon)$ and $\varepsilon = 0.5$, and a budget equal to 25% of the total welfare loss from the flood. Transfers proportional to own flooding reduce the social welfare loss by 37%. This exceeds 25% because g is convex and these transfers go mostly to households




Figure 11: Budget saved by optimal targeting, by budget and by curvature of the planner's objective

Notes: Each point is the share of the budget a rule saves against transfers proportional to own flooding while delivering the same social welfare, from the planner problem of equation 22 with $g(\ell) = \ell^{1+\epsilon}/(1+\epsilon)$. A: the budget varies, as a share of the sum of households' welfare losses, with $\epsilon = 0.5$. B: ϵ varies, with the budget at 25 percent. Dashed lines mark the paper's values.

with large losses. Choosing the schedule optimally, but still conditioning only on own flooding, achieves the same reduction with 97% of the budget. Conditioning transfers on predicted welfare losses achieves it with 82%. Almost all of the gain therefore comes from conditioning transfers on predicted losses, not from using them to design the schedule.

Figure 11 shows how these savings vary with the budget and with ϵ . Savings are smallest at a budget of around 25% of the total loss, and larger at both smaller and larger budgets (Panel A). At small budgets, both optimal allocations gain by concentrating transfers on the largest losses. At large budgets, targeting on predicted losses also avoids overpaying the most flooded hexagons and reaches those whose losses are due to flooding elsewhere. The gap between the two allocations, the value of the predicted losses, is 12 to 15 percentage points at every budget. This value is zero when g is linear and rises with the planner's aversion to large losses, to 30% at $\epsilon = 1$ and 62% at $\epsilon = 3$ (Panel B).

8.2 Other floods

Our analysis so far has focused on a single case study, a large and spatially dispersed flood in a low-income economy. To extrapolate from our findings to other disasters and economies, and to draw broader lessons for the design of disaster assistance, we use the model to study how our results depend on features of the flood and of the economy. We focus on three: how integrated the economy was before the flood, the spatial concentration of the flood, and its size.

8.2.1 Pre-flood economic integration

We have argued that the incidence of a flood differs from its physical footprint because locations are linked through trade and commuting, and these economic linkages share the costs of a disaster across space. It follows that the costs of trade and commuting are an important determinant of a flood's incidence.

To see why, consider an economy in which these costs are high, so that residents mostly work and shop where they live (π_{ii} and π_{ii}^G close to one). These shares also barely respond to wages and prices, since the elasticities of π_{ii} with respect to the local wage and of π_{ii}^G with respect to the local price are $\theta_L(1 - \pi_{ii})$ and $-\theta_T(1 - \pi_{ii}^G)$. Three implications follow. First, by equation 17, residents of i are exposed to first order only to changes in wages and prices in i . Second, flooding in other locations affects wages and prices in i only through changes in labour supply to i or in demand for its goods. Since commuting and trade shares are unresponsive, both are approximately unchanged, and wages and prices in i depend only on flooding in i . Third, residents of i cannot substitute away from these changes by working or buying elsewhere. The welfare loss from a flood is therefore borne by the locations it floods.

By contrast, when trade and commuting costs are low, residents work and shop widely, so π_{ii} and π_{ii}^G are small and highly responsive to wages and prices. Residents of i are then exposed to wages and prices in many locations, each of which depends on flooding everywhere, because workers and buyers reallocate towards wherever wages are higher and prices lower. The productivity loss from a flood is therefore shared across the whole economy.

8.2.2 Spatial dispersion of the flood

Holding the share of locations flooded fixed, the spatial pattern of a flood changes who bears its cost. Since π_{ij} and π_{ij}^G fall with distance, residents of i work and shop mainly in locations close to i , and the spillovers from flooding in i fall mainly on those locations. In a concentrated flood, flooded locations tend to have flooded neighbours and unflooded locations unflooded neighbours. To first order, flooded locations are therefore exposed to flooding both directly and indirectly, while unflooded locations are exposed to neither. Moreover, residents of flooded locations cannot offset these changes by shifting their labour supply or goods demand elsewhere, because the locations they can reach are flooded too.

In a dispersed flood, the locations close to a flooded location are dry. Locations tend to be exposed either directly or indirectly, but not both. Moreover, residents of flooded locations can reach unflooded locations in which to supply labour and buy goods, and so share their loss with them. It follows that own flooding is a good predictor of welfare loss in a concentrated flood, when direct and indirect exposure are strongly correlated, and a worse one in a dispersed flood.

We consider counterfactual floods that keep the same distribution of flooding across hexagons as in 2022, but rearrange it so that flooding is more or less spatially concentrated. In the most dispersed flood, unflooded hexagons bear 30% of the total welfare loss, compared with 24% in the most concentrated. The rank R^2 between own flooding and welfare loss rises only slightly, from 0.67 to 0.71.

8.2.3 Size of the flood

In this model, the aggregate cost of a flood is increasing in the number of locations flooded, holding the intensity of flooding fixed. To first order, by equation 17, the aggregate loss is linear in the number of locations flooded, since trade and commuting redistribute the loss across locations without changing its total. Taking into account second order effects, residents of flooded locations offset part of their loss by shifting labour supply and goods demand towards unflooded locations, and their ability to do so is greatest when few locations are flooded. The aggregate cost is therefore mildly

convex, and more so the lower are trade and commuting costs and the higher are θ_L and θ_T . In autarky, there is no substitution, and it is linear.

The size of a flood also changes its incidence. Linearising around the pre-flood equilibrium, we can write the change in welfare in the form

$$d \ln V_i^w = \beta d \ln B_i + \sum_k F_{ik} a_k, \quad (25)$$

where F_{ik} is the general equilibrium effect of a productivity shock in k on welfare in i , which depends on the pre-flood trade and commuting shares and on the model's parameters, and $\sum_k F_{ik} = 1 - \beta$.

Consider a flood that lowers productivity by a and housing by b in a set of locations $\mathcal{F}$. The welfare loss in i , to first order, is

$$-d \ln V_i^w = \beta b \mathbf{1}\{i \in \mathcal{F}\} + a \sum_{k \in \mathcal{F}} F_{ik}. \quad (26)$$

The first term is the loss due to housing destruction, which is the same (βb) for every flooded location regardless of the scale of the flood. The second is the loss due to productivity impacts, which is increasing in the number of flooded locations, whether or not i is itself flooded. As the flood grows, the second term rises in every location while the first does not. The ratio of the loss of a resident of a flooded location to that of a resident of an unflooded location therefore falls. It falls more slowly in a concentrated flood, since newly flooded locations then border locations that are already flooded, and so the productivity impacts fall mainly on locations that are already flooded.

The implication is that own flooding predicts welfare loss less well as the flood grows, up to a point. Own flooding determines the first term in equation 26, so what matters is how much the second term varies across locations. It is zero in every location when none is flooded, and equal to $a(1 - \beta)$ in every location when all are flooded, so it varies most at intermediate sizes. Own flooding therefore ranks welfare losses well in small floods and in very large floods, and worst in between.

Consider counterfactual floods that vary the share of hexagons flooded, holding the distribution of flooding among flooded hexagons and spatial concentration fixed at their 2022 values. The welfare loss per resident of a flooded hexagon is 7.1%, 8.4%, 9.9% and 11.3% when 5%, 20%, 40% and 60%

of hexagons are flooded, and per resident of an unflooded hexagon 1.9%, 3.1%, 4.5% and 5.5%, so the ratio falls from 3.8 to 2.1. The rank R^2 between own flooding and welfare loss is 0.79 for the smallest floods, falls to 0.66 when 30% of hexagons are flooded, and rises slightly thereafter.

8.2.4 Summary

Own flooding is therefore a worse guide to welfare loss when floods are large and dispersed and when trade and commuting costs are neither very high nor very low. Many disasters are much smaller than the 2022 floods, and our results suggest that when a disaster affects few locations, physical exposure is a good guide to welfare loss. However, large floods are projected to become more frequent with climate change, meaning that accounting for indirect exposure will become increasingly important. Our results on trade and commuting costs also suggest that the spillovers we document are not specific to a low-income setting. Trade and commuting costs are typically lower in high-income economies, and our model would predict that a larger share of a flood's cost is then borne outside the flooded locations.

8.3 General-equilibrium effects of aid

The planner's problem above treats aid as a lump-sum transfer of utility, which does not change household behaviour and so has no general-equilibrium effects. Aid that restores productive capacity in the locations where it is spent does have such effects, and the optimal allocation then depends on the spillovers from that spending.

We have so far considered cash transfers. However it is possible that disaster assistance could restore productive capacity, by draining fields, repairing irrigation or replacing livestock and machinery. Aid of this kind would reverse (at least partially) the productivity loss in the location it is spent, and so reverses the losses that location imposed on others.

We model this kind of aid as spending y_j in location j that raises agricultural productivity T_j^A . The planner solves equation 22 with y_j in place of utility transfers. We assume for now that they cannot increase productivity in a location above its pre-flood level ($T_j^A(f_j, y_j) \leq \bar{T}_j^A$).

The first-order condition, equation 23, becomes

$$\frac{\partial \ln T_j^A}{\partial y_j} \sum_i \omega_i g'(\ell_i) \frac{\partial V_i}{\partial \ln T_j^A} = \mu - \lambda_j + v_j, \quad v_j \geq 0, \quad v_j (\bar{T}_j^A - T_j^A) = 0, \quad (27)$$

where $\ell_i = V_i(\mathbf{0}, \mathbf{0}) - V_i(\mathbf{f}, \mathbf{y})$ is the remaining loss in i and v_j is the multiplier on the cap.

The left-hand side of equation 27 is the social value of a rupee spent in j . An extra rupee in j raises productivity there by $\partial \ln T_j^A / \partial y_j$. This raises welfare in location i by $\partial V_i / \partial \ln T_j^A$ per unit of productivity. The planner values an increase in V_i more highly if i has lost more given the flood and the aid so far, as measured by $\omega_i g'(\ell_i)$.

The planner spends more in j if more households are exposed to its productivity (for example, j is a large employer or food supplier), or if those households have lost more as a result of the flood (for example, they suffered direct damage or sit in an area of widespread flooding). Intuitively, the value of aid in j is its contribution to the market access of all locations, weighted by the losses suffered by residents of these locations. Aid therefore does not necessarily flow to the most damaged locations, since restoring productivity in the locations they trade with and commute to can reduce their losses by more.

Given the constraint that aid can only restore productivity to pre-flood levels, it is allocated only to flooded locations. Relaxing this constraint, the planner may direct some aid to locations that were not flooded but are highly connected to flooded regions.

9 CONCLUSION

This paper shows that the economic incidence of a natural disaster can differ substantially from its physical footprint. Using original surveys of 5,100 rural households after Pakistan's 2022 floods and a recentered instrument based on a hydrological model, we find that flooding in the surrounding area lowers household consumption holding flooding at home fixed. Flooding more than 5km from the home accounts for nearly two-thirds of the effect of flooding within 15km. The same flooding reduces wage employment and raises food prices, consistent with spillovers through labour and product markets.

To quantify incidence in general equilibrium, we build a spatial model of trade and commuting. For the 2022 floods in Sindh, the model implies that welfare fell in every location, although two in five experienced almost no flooding, and that 73% of the average loss was due to flooding elsewhere. Targeting aid on predicted losses rather than own flooding would have achieved the same welfare at 18% lower cost. The gap between physical exposure and incidence is largest for large, dispersed floods.

The broader lesson is that trade and commuting with unflooded areas insure households against localised shocks, and that this insurance breaks down when the shock is common to 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, ignoring the divergence between incidence and physical footprint will increasingly lead us to underestimate the impact of disasters and to misallocate aid.

A GEOMETRY OF THE EXPOSURE MEASURES

Table A.1: Area and within-district variation of the exposure measures

		Within-district p10 to p90 gap		
		Households	Housing hexes	Cropland hexes
<i>Panel A: flooding within the radius</i>				
Radius (km)	Area (km ²)			
0.5	0.8	0.96	0.99	1.08
1	3.1	0.95	0.95	1.01
2	12.6	0.96	0.89	1.04
3	28.3	0.92	0.82	1.00
4	50.3	0.91	0.83	0.89
5	78.5	0.82	0.77	0.78
6	113.1	0.80	0.74	0.75
7	153.9	0.80	0.71	0.72
8	201.1	0.76	0.68	0.69
9	254.5	0.66	0.61	0.63
10	314.2	0.59	0.55	0.61
11	380.1	0.56	0.52	0.58
12	452.4	0.52	0.49	0.55
13	530.9	0.49	0.45	0.51
14	615.8	0.46	0.40	0.47
15	706.9	0.42	0.36	0.43
<i>Panel B: own flooding within the cutoff / nearby flooding to 15 km</i>				
Cutoff (km)	Own share (%)			
0.5	0.1	0.96 / 0.42	0.99 / 0.36	1.08 / 0.43
1	0.4	0.95 / 0.43	0.95 / 0.36	1.01 / 0.43
2	1.8	0.96 / 0.43	0.89 / 0.36	1.04 / 0.44
3	4.0	0.92 / 0.44	0.82 / 0.36	1.00 / 0.44
4	7.1	0.91 / 0.44	0.83 / 0.36	0.89 / 0.44
5	11.1	0.82 / 0.43	0.77 / 0.37	0.78 / 0.44
6	16.0	0.80 / 0.42	0.74 / 0.36	0.75 / 0.43
7	21.8	0.80 / 0.40	0.71 / 0.36	0.72 / 0.42
8	28.4	0.76 / 0.38	0.68 / 0.35	0.69 / 0.39
9	36.0	0.66 / 0.37	0.61 / 0.35	0.63 / 0.38
10	44.4	0.59 / 0.36	0.55 / 0.34	0.61 / 0.38
<i>Panel C: move of own / nearby flooding along the 15 km p10 to p90, households</i>				
Cutoff (km)				
0.5		0.28 / 0.42		
1		0.30 / 0.42		
2		0.34 / 0.42		
3		0.37 / 0.42		
4		0.40 / 0.42		
5		0.44 / 0.42		
6		0.46 / 0.41		
7		0.48 / 0.41		
8		0.49 / 0.39		
9		0.50 / 0.38		
10		0.51 / 0.36		

Notes: Gap: 10th to 90th percentile range of the regressor net of district means, the unit the coefficients are scaled to. Housing hexes have at least 50 buildings. Own share: area within the cutoff over the 15 km disk.

Table A.2: Own and nearby coefficients per unit of the flood index, by own cutoff

Own cutoff (km)	Real expenditure (%)	Work for others (pp)	Commutes >1 km (pp)	Food prices (%)	Has loans (pp)	Sold assets (pp)	Building damage (pp)	Crop greenness (% normal)	Expenditure, no relief (%)	Expenditure, relief (%)
<i>Own flooding, within the cutoff</i>										
0.5	-1.50 (5.43)	0.86 (1.66)	7.42 (1.55)	0.30 (1.97)	-0.18 (1.57)	2.02 (1.72)	5.93 (2.80)	-30.23 (4.52)	-15.09 (13.53)	1.44 (5.78)
1	-3.80 (4.31)	1.81 (1.60)	7.25 (1.49)	1.89 (1.64)	-1.52 (1.52)	0.98 (1.65)	6.20 (2.78)	-32.77 (4.88)	-11.91 (9.28)	-3.83 (4.57)
2	-4.68 (3.37)	2.18 (1.65)	7.44 (1.85)	2.17 (1.91)	-1.86 (1.94)	0.69 (1.84)	5.73 (2.89)	-34.85 (5.37)	-8.32 (5.95)	-6.07 (4.74)
3	-5.92 (3.43)	2.65 (2.15)	8.33 (2.31)	2.04 (2.55)	-3.17 (2.48)	0.70 (2.47)	6.24 (3.32)	-36.27 (6.20)	-10.42 (6.43)	-5.41 (6.14)
4	-7.71 (4.10)	4.08 (2.79)	9.65 (3.01)	1.91 (2.98)	-4.01 (2.96)	0.85 (3.14)	6.70 (3.97)	-36.31 (7.59)	-13.65 (8.05)	-5.33 (7.69)
5	-8.12 (4.89)	5.14 (3.62)	11.08 (3.79)	1.13 (3.43)	-4.42 (3.31)	1.05 (3.66)	7.20 (4.68)	-34.90 (9.18)	-14.72 (9.61)	-3.71 (9.22)
<i>Nearby flooding, cutoff to 15 km</i>										
0.5	-22.93 (9.80)	-21.28 (8.37)	-5.99 (4.97)	19.00 (7.70)	-11.49 (5.72)	5.17 (2.90)	-0.09 (4.49)	-5.31 (3.90)	-33.06 (13.82)	-20.09 (10.99)
1	-21.24 (9.83)	-21.91 (8.22)	-6.36 (5.14)	17.88 (7.99)	-10.56 (5.87)	5.74 (2.91)	-0.08 (4.65)	-3.25 (4.07)	-33.28 (13.92)	-14.90 (11.47)
2	-20.17 (9.60)	-22.24 (7.93)	-7.06 (5.37)	17.46 (8.38)	-10.15 (6.13)	5.86 (2.83)	-0.45 (4.77)	0.23 (4.44)	-33.55 (12.92)	-12.50 (11.65)
3	-18.68 (9.20)	-22.77 (7.93)	-8.44 (5.22)	17.34 (8.84)	-8.87 (6.42)	5.79 (2.81)	-1.73 (5.29)	3.42 (4.81)	-31.14 (11.94)	-12.89 (12.26)
4	-16.44 (9.22)	-24.41 (8.18)	-10.56 (5.06)	17.31 (9.21)	-7.77 (6.87)	5.59 (3.08)	-4.08 (6.29)	6.79 (5.81)	-27.39 (11.51)	-12.58 (12.97)
5	-15.43 (9.28)	-26.12 (8.90)	-13.07 (5.23)	18.12 (9.70)	-7.03 (7.35)	5.36 (3.54)	-6.46 (7.73)	9.77 (7.45)	-25.20 (11.20)	-14.17 (13.65)

Notes: IV estimates, district fixed effects; one regression per cutoff on own flooding within the cutoff and nearby flooding from the cutoff to 15 km. Coefficients per unit of the flood index, times 100; Bartlett-kernel Conley standard errors in parentheses. Samples as in the corresponding figures. The moves of each regressor along the 15 km p10 to p90 are in Table A.1.

B ADDITIONAL FIGURES




Figure 12: Effect of flooding on real expenditure, households reporting relief

Notes: IV estimates, district fixed effects, households reporting relief: coefficient plus interaction from one regression per point on the pooled sample with flooding, a relief indicator and their interaction, instruments the recentred prediction and its interaction. Outcome: log real expenditure per capita. A: flooding within the radius, per within-district p10 to p90 of the regressor, times 100. B: the 15 km effect split into the part through own flooding within the cutoff (blue) and through nearby flooding from the cutoff to 15 km (orange), each coefficient times the move in that regressor along the 15 km p10 to p90, so the two add up to the dashed line, the 15 km total from A. Per-unit coefficients in Table A.2. 90 percent Conley intervals, Bartlett kernel. N = 4,178 households.




Figure 13: Effect of flooding on commuting more than 1km

Notes: IV estimates, district fixed effects, all households. Outcome: any man in the household commutes more than 1 km, $N = 4,478$ households. A: flooding within the radius, per within-district p10 to p90 of the regressor, times 100. B: the 15 km effect split into the part through own flooding within the cutoff (blue) and through nearby flooding from the cutoff to 15 km (orange), each coefficient times the move in that regressor along the 15 km p10 to p90, so the two add up to the dashed line, the 15 km total from A. Per-unit coefficients in Table A.2. 90 percent Conley intervals, Bartlett kernel.




Figure 14: Effect of flooding on employment and self-employment

Notes: IV estimates, district fixed effects, all households. Top: share of the household's men in any employment, $N = 4,478$ households. Bottom: share of the household's men self-employed, $N = 4,478$ households. A: flooding within the radius, per within-district p10 to p90 of the regressor, times 100. B: the 15 km effect split into the part through own flooding within the cutoff (blue) and through nearby flooding from the cutoff to 15 km (orange), each coefficient times the move in that regressor along the 15 km p10 to p90, so the two add up to the dashed line, the 15 km total from A. Per-unit coefficients in Table A.2. 90 percent Conley intervals, Bartlett kernel.

B.1 Wage employment by pre-flood sector




Figure 15: Effect of flooding on working for others, households whose men were all in agriculture before the floods

Notes: IV estimates, district fixed effects, households whose employed men were all in agriculture before the floods. Outcome: share of the household's men working for others, $N = 1,830$ households. A: flooding within the radius, per within-district p10 to p90 of the regressor, times 100. B: the 15 km effect split into the part through own flooding within the cutoff (blue) and through nearby flooding from the cutoff to 15 km (orange), each coefficient times the move in that regressor along the 15 km p10 to p90, so the two add up to the dashed line, the 15 km total from A. Per-unit coefficients in Table A.2. 90 percent Conley intervals, Bartlett kernel.




Figure 16: Effect of flooding on working for others, households with at least one man outside agriculture before the floods

Notes: IV estimates, district fixed effects, households with at least one man outside agriculture before the floods. Outcome: share of the household's men working for others, $N = 2,648$ households. A: flooding within the radius, per within-district p10 to p90 of the regressor, times 100. B: the 15 km effect split into the part through own flooding within the cutoff (blue) and through nearby flooding from the cutoff to 15 km (orange), each coefficient times the move in that regressor along the 15 km p10 to p90, so the two add up to the dashed line, the 15 km total from A. Per-unit coefficients in Table A.2. 90 percent Conley intervals, Bartlett kernel.

C PARAMETER ESTIMATION

Table C.1: Gravity in road travel time, by trade year

	2018/19	2019/20	2020/21	2021/22	2022/23	2023/24	2024/25	All
<i>Panel A. OLS, log value</i>								
Log travel time	-1.51*** (0.240)	-1.38*** (0.198)	-1.40*** (0.164)	-1.19*** (0.176)	-1.27*** (0.157)	-1.47*** (0.175)	-1.16*** (0.159)	-1.32*** (0.067)
Observations	611	866	950	999	1,202	1,081	1,220	6,929
<i>Panel B. PPML, value in levels</i>								
Log travel time	-2.51*** (0.254)	-2.06*** (0.234)	-1.82*** (0.193)	-1.71*** (0.205)	-1.84*** (0.210)	-2.15*** (0.213)	-2.05*** (0.174)	-1.99*** (0.081)
Observations	5,616	6,623	7,379	6,692	8,711	8,189	9,198	52,408
<i>Panel C. PPML on the OLS sample</i>								
Log travel time	-1.61*** (0.312)	-1.50*** (0.258)	-1.32*** (0.209)	-1.33*** (0.217)	-1.55*** (0.250)	-1.76*** (0.246)	-1.70*** (0.192)	-1.55*** (0.092)
Observations	611	866	950	999	1,202	1,081	1,220	6,929

Notes: Distance is road travel time in minutes from the least-time routes over the speed-assigned network; internal trade keeps the Head-Mayer internal distance at 15 km/h. All crops summed into one flow per origin-destination-year. Heteroskedasticity-robust standard errors beneath each estimate. Trade years run 1 May to 30 April. The year columns have origin and destination fixed effects; the pooled column has origin $\times$ year and destination $\times$ year fixed effects. OLS uses only pairs that traded, because the outcome is in logs; Panel B also uses pairs with no recorded flow, which is why its samples are larger. Panel C drops those pairs again, so it differs from Panel A only in the estimator and from Panel B only in the sample. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

REFERENCES

- Adão, Rodrigo, Costas Arkolakis, and Federico Esposito**, “Spatial Linkages, Global Shocks, and Local Labor Markets: Theory and Evidence,” Working Paper 25544, National Bureau of Economic Research 2019. [Return to page [28](#)]
- Ahlfeldt, Gabriel M., Stephen J. Redding, Daniel M. Sturm, and Nikolaus Wolf**, “The Economics of Density: Evidence from the Berlin Wall,” *Econometrica*, 2015, 83 (6), 2127–2189. [Referenced on pages [3](#) and [28](#)]
- 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 [5](#)]
- 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 [5](#)]
- Allen, Treb and David Atkin**, “Volatility and the Gains from Trade,” *Econometrica*, 2022, 90 (5), 2053–2092. [Referenced on pages [5](#) and [26](#)]
- Associated Press of Pakistan**, “BISP Disburses over 97 Percent of Flood Relief Cash Assistance among Flood Affected Families,” 30 October 2022. <https://www.app.com.pk/national/bisp-disburses-over-97-percent-of-flood-relief-cash-assistance-among-flood-affected-families/>. [Return to page [39](#)]
- Athey, Susan and Stefan Wager**, “Policy Learning with Observational Data,” *Econometrica*, 2021, 89 (1), 133–161. [Return to page [5](#)]
- 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 [5](#)]
- 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 [7](#)]

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 [5](#)]

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 [5](#)]

Bates, Paul D., Matthew S. Horritt, and Timothy J. Fewtrell, “A simple inertial formulation of the shallow water equations for efficient two-dimensional flood inundation modelling,” *Journal of Hydrology*, 2010, 387 (1–2), 33–45. [Referenced on pages [11](#) and [16](#)]

Betterle, Andrea and Peter Salamon, “Water depth estimate and flood extent enhancement for satellite-based inundation maps,” *Natural Hazards and Earth System Sciences*, 2024, 24 (8), 2817–2836. [Return to page [10](#)]

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 [5](#)]

Borusyak, Kirill and Peter Hull, “Nonrandom Exposure to Exogenous Shocks,” *Econometrica*, 2023, 91 (6), 2155–2185. [Referenced on pages [1](#), [12](#), and [13](#)]

— and —, “Optimal Formula Instruments,” Working Paper 33594, National Bureau of Economic Research 2025. [Referenced on pages [14](#) and [16](#)]

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 [5](#)]

Breza, Emily, Supreet Kaur, and Yogita Shamdasani, “Labor Rationing,” *American Economic Review*, 2021, 111 (10), 3184–3224. [Return to page [25](#)]

Brooks, Wyatt and Kevin Donovan, “Eliminating Uncertainty in Market Access: The Impact of New Bridges in Rural Nicaragua,” *Econometrica*, 2020, 88 (5), 1965–1997. [Referenced on pages [5](#) and [24](#)]

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 [7](#)]

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 [5](#)]

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 [5](#)]

Castro-Vincenzi, Juan Manuel, “Climate Hazards and Resilience in the Global Car Industry,” Working Paper, Princeton University 2024. [Return to page [5](#)]

Chamberlain, Gary, “Asymptotic efficiency in estimation with conditional moment restrictions,” *Journal of Econometrics*, 1987, 34 (3), 305–334. [Return to page [16](#)]

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 [16](#)]

Chodorow-Reich, Gabriel, “Regional Data in Macroeconomics: Some Advice for Practitioners,” *Journal of Economic Dynamics and Control*, 2020, 115, 103875. [Return to page [28](#)]

Conley, Timothy G., “GMM Estimation with Cross Sectional Dependence,” *Journal of Econometrics*, 1999, 92 (1), 1–45. [Return to page [16](#)]

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. [Return to page [5](#)]

- Donald, Eric, Masao Fukui, and Yuhei Miyauchi**, “Unpacking Aggregate Welfare in a Spatial Economy,” *Review of Economic Studies*, 2026. Forthcoming. [Return to page [5](#)]
- Eaton, Jonathan and Samuel Kortum**, “Technology, Geography, and Trade,” *Econometrica*, 2002, 70 (5), 1741–1779. [Referenced on pages [3](#) and [28](#)]
- Elbers, Chris, Tomoki Fujii, Peter Lanjouw, Berk Özler, and Wesley Yin**, “Poverty Alleviation through Geographic Targeting: How Much Does Disaggregation Help?,” *Journal of Development Economics*, 2007, 83 (1), 198–213. [Return to page [5](#)]
- Fajgelbaum, Pablo D. and Cecile Gaubert**, “Optimal Spatial Policies, Geography, and Sorting,” *Quarterly Journal of Economics*, 2020, 135 (2), 959–1036. [Return to page [5](#)]
- Funk, Chris, Pete Peterson, Martin Landsfeld, Diego Pedreros, James Verdin, Shrad-dhanand 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 [6](#)]
- 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 [5](#)]
- 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 [5](#)]
- 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 [5](#)]

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 [5](#)]

Haushofer, Johannes, Paul Niehaus, Carlos Paramo, Edward Miguel, and Michael Walker, “Targeting Impact versus Deprivation,” *American Economic Review*, 2025, 115 (6), 1936–1974. [Return to page [5](#)]

Huber, Kilian, “Estimating General Equilibrium Spillovers of Large-Scale Shocks,” *The Review of Financial Studies*, 2023, 36 (4), 1548–1584. [Return to page [20](#)]

Kaur, Supreet, “Nominal Wage Rigidity in Village Labor Markets,” *American Economic Review*, 2019, 109 (10), 3585–3616. [Return to page [25](#)]

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 [7](#)]

Kirchberger, Martina, “Natural Disasters and Labor Markets,” *Journal of Development Economics*, 2017, 125, 40–58. [Return to page [5](#)]

Kitagawa, Toru and Aleksey Tetenov, “Who Should Be Treated? Empirical Welfare Maximization Methods for Treatment Choice,” *Econometrica*, 2018, 86 (2), 591–616. [Return to page [5](#)]

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 [5](#)]

Ministry of Finance, Government of Pakistan, “Social Protection,” in “Pakistan Economic Survey 2022–23” 2023, chapter 16. https://www.finance.gov.pk/survey/chapters_23/16_Social_Protection.pdf. [Return to page [39](#)]

Morten, Melanie, “Temporary Migration and Endogenous Risk Sharing in Village India,” *Journal of Political Economy*, 2019, 127 (1), 1–46. [Return to page [5](#)]

Munshi, Kaivan and Mark Rosenzweig, “Networks and Misallocation: Insurance, Migration, and the Rural-Urban Wage Gap,” *American Economic Review*, 2016, *106* (1), 46–98. [Return to page [5](#)]

Nakamura, Emi and Jón Steinsson, “Fiscal Stimulus in a Monetary Union: Evidence from US Regions,” *American Economic Review*, 2014, *104* (3), 753–792. [Return to page [28](#)]

Neal, Jeffrey, Guy Schumann, and Paul Bates, “A subgrid channel model for simulating river hydraulics and floodplain inundation over large and data sparse areas,” *Water Resources Research*, 2012, *48* (11), W11506. [Referenced on pages [11](#) and [16](#)]

Newbery, David M. G. and Joseph E. Stiglitz, “Pareto Inferior Trade,” *Review of Economic Studies*, 1984, *51* (1), 1–12. [Return to page [4](#)]

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 [6](#)]

Patel, Dev, “Floods,” 2024. Working paper. [Referenced on pages [5](#), [7](#), and [10](#)]

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 [7](#)]

Ravallion, Martin, “Poverty Alleviation through Regional Targeting: A Case Study for Indonesia,” in Karla Hoff, Avishay Braverman, and Joseph E. Stiglitz, eds., *The Economics of Rural Organization: Theory, Practice, and Policy*, Oxford: Oxford University Press, 1993. [Return to page [5](#)]

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 [7](#)]

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 7]

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

Sturm, Daniel M., Kohei Takeda, and Anthony J. Venables, “How Useful are Quantitative Urban Models for Cities in Developing Countries? Evidence from Dhaka,” Policy Research Working Paper 9910, World Bank 2021. [Return to page 34]

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 6]

Townsend, Robert M., “Risk and Insurance in Village India,” *Econometrica*, 1994, 62 (3), 539–591. [Return to page 5]

Webster, Elizabeth M., “FEMA’s Individuals and Households Program (IHP): Implementation and Considerations for Congress,” Report R47015, Congressional Research Service 2023. Updated 30 June 2023. <https://crsreports.congress.gov/product/pdf/R/R47015>. [Return to page 42]

Wolf, Christian K., “The Missing Intercept: A Demand Equivalence Approach,” *American Economic Review*, 2023, 113 (8), 2232–2269. [Return to page 28]

World Bank, GFDRR, European Union, and United Nations, “Post-Disaster Needs Assessments, Volume A: Guidelines,” Technical Report, World Bank, Global Facility for Disaster Reduction and Recovery (GFDRR), European Union and United Nations 2013. [Return to page 42]

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-pakistan/> 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 6]