Vulnerable Markets: Impact of Extreme Flooding on Agriculture Supply Networks*

Sher Afghan Asad† **Omar Hayat Gondal**‡ **Farah Said**§

This paper studies how extreme climate shocks propagate through agricultural supply networks, using the 2022 floods in Pakistan as a natural experiment. We assemble novel administrative data on agricultural commodity arrivals at 36 regulated wholesale markets, tracing flows from 223 source regions, and link these records to high-resolution satellite measures of flood exposure constructed for this study. Our empirical strategy combines a staggered event-study design at production origins with a shift-share exposure framework at destination markets, allowing us to separately identify upstream production shocks, downstream market impacts, and network adjustment. We document three main results. First, flooding causes large and persistent supply losses: arrivals from flooded origins decrease by roughly 50 percent, and downstream market arrivals decline by up to 65 percent, particularly for perishable crops. Second, price adjustment is uneven across the supply chain: retail prices rise quickly and remain elevated, while wholesale prices respond more gradually, implying temporary increases in intermediary margins. Third, network structure shapes the incidence of the shock: the downstream supply collapse is disproportionately driven by disruptions to long-distance sourcing relationships, indicating that spatially extended supply links are especially vulnerable during large-scale floods. We further show that historically flood-prone origins exhibit anticipatory behavior and faster recovery, consistent with differential preparedness.

*We thank Robin Burgess, Shan He, Amen Jalal, Muhammad Bin Khalid, participants at the Econometric Society Interdisciplinary Frontiers - Economics+Climate Science (ESIFCLIM) conference 2025 in Barcelona, Western Economic Association Conference 2024 in Seattle, Applied Development Economics Conference 2024 in Lahore, and Pathways to Development Conference 2023 in Lahore for their helpful comments and suggestions. We also thank Somia Asim, Sameer Hayee, Azka Shahid, Zoha Awais, Massab Qayum, Danish Sarwar, and Waqas Ahmed for their excellent research assistance. The authors acknowledge the financial support received from the International Growth Center (IGC) (Grant PAK-23156), and the Lahore University of Management Sciences (LUMS) - Research Initiative for Societal Challenges (Grant 004).

†Lahore University of Management Sciences (sherafghan@lums.edu.pk)

‡Washington University in St. Louis (omargondal@wustl.edu)

§Lahore University of Management Sciences (farah_said@lums.edu.pk)

1 Introduction

Climate shocks are no longer rare tail events—they are recurring and causing systemic disruptions with profound consequences for food systems and market stability (Hsiang and Kopp, 2018). Yet we still know relatively little about how these shocks propagate through agricultural supply chains in low-income settings. Most research has focused on direct, farm-level impacts—on yields, labor, or long-run adaptation, while overlooking how production shocks translate into downstream market outcomes. This matters acutely in the Global South, where fragmented supply chains, weak infrastructure, and tenancy-related rigidities make food systems especially vulnerable (Douglas, 2009; Hanna and Oliva, 2016), and where the state plays an active role in managing staple commodity flows.

As climate volatility continues to destabilize food and input markets, understanding these downstream disruptions is essential for designing timely and spatially targeted policy responses. Our analysis is motivated by both the urgency of these shocks and the absence of scalable tools to monitor their propagation through informal and regulated systems.

Pakistan, ranked among the countries most vulnerable to climate shocks (Eckstein et al., 2021), has faced repeated extreme weather events with devastating consequences. In 2022, monsoon floods submerged nearly one-third of the country, affected over 30 million people, and inflicted damages exceeding $30 billion. The disaster struck during the *kharif* harvest, inundating farms, destroying infrastructure, and revealing deep fragilities in agricultural logistics. These conditions, combined with the country's reliance on perishable crops and government-mediated commodity flows—make Pakistan, and especially the populous province of Punjab, a critical test case for understanding how climate shocks disrupt food supply chains.

This paper studies how the 2022 floods disrupted the supply of key agricultural commodities into Punjab, Pakistan's most populous province and largest consumption hub. We leverage a unique high-frequency administrative dataset that records the daily arrival of agricultural goods into 128 major wholesale markets (*mandis*) in Punjab. Using georeferenced source locations for each shipment and satellite-based flood maps, we trace how supply networks linking flood-affected source regions to destination districts responded to inundation. Our empirical strategy

uses a staggered event study design that exploits variation in both the timing and intensity of flood exposure across space. This framework allows us to separately estimate anticipatory responses before inundation, as well as post-flood disruptions and the pace of recovery. To strengthen causal identification, we incorporate spatial flood risk profiles, crop-level heterogeneity, and road network routing algorithms that proxy for real-world transport links between source and destination. Our approach allows us to analyze the propagation of climate shocks across fragmented agricultural supply chains using publicly available geospatial tools—without relying on proprietary tracking or transaction data.

We document four key findings. First, agricultural supply chains experienced severe disruptions following the 2022 floods, with quantity supplied falling by up to 50% at the peak of the crisis, and downstream market arrivals falling by as much as 65%. Second, the impacts varied sharply by crop type: perishable commodities - fruits and vegetables - saw the largest and most prolonged disruptions, while grain flows remained stable, likely due to greater storability, larger buffer stocks, and public procurement programs. We also find evidence of short-run anticipatory behavior: in flood-prone source regions, shipments temporarily increased before inundation, suggesting producers and intermediaries accelerated deliveries in expectation of supply disruptions. Fourth, flood shocks triggered delayed but substantial prices increases throughout the supply chain, with wholesale and retail prices rising by 49% and 65%, respectively for vegetables, and retailer margins widening by up to 27%. Finally, we find that government enforcement of price ceilings responded gradually but meaningfully, with ceiling adjustments exceeding 120% in the exposed markets, coinciding with a moderation in margins, highlighting the role of state regulation in shaping distributional outcomes during climate-induced crises.

Our study makes several contributions to the literature. First, we move beyond the dominant focus on production-side outcomes in climate-agriculture research (e.g., Deschênes and Greenstone, 2012; Dell et al., 2012; Burgess et al., 2017), and instead document market-level disruptions in agricultural supply chains—an underexplored but crucial link with important implications for food security.

Second, we contribute to the literature on climate shocks and supply chains, which has largely focused on industrial or high-income settings (Cavallo et al., 2014; Barrot and Sauvagnat, 2016;

Boehm et al., 2019). While recent work has explored firm- or transport-level resilience to flooding in South Asia (e.g., Balboni et al., 2024), we focus on the intermediary space of agricultural markets, where physical flows, pricing, and state oversight converge. Our study provides new high-frequency evidence from a developing-country agricultural context, showing how climate shocks propagate spatially and temporally through spatially fragmented and infrastructure-constrained supply networks.

Third, we present novel evidence of short-run anticipatory behavior in agriculture: producers and intermediaries in flood-prone areas accelerate shipments ahead of expected inundation. This form of preemptive adjustment is rarely documented compared to the longer-run adaptation literature. By leveraging fine-grained flood risk variation, we isolate this anticipatory response from general seasonality or pre-trends.

Fourth, we offer new insights into the role of state regulation in crisis-era market dynamics. Existing literature offers mixed evidence on the effectiveness of crisis-era interventions such as price controls and public procurement (Timmer, 2010; Headey, 2011). We show that price spikes and margin increases were sharper in perishable goods but eventually moderated in markets where price ceiling enforcement was stronger. This suggests that even in the absence of large-scale buffer stock releases or public procurement, regulatory enforcement can play a stabilizing role by curbing intermediary markups during periods of acute scarcity.

Classic work shows that lower trade costs improve price integration and reduce dispersion (Atkin and Donaldson, 2015; Donaldson, 2018; Allen and Arkolakis, 2014), and that information and competition shape food-market performance (Jensen, 2007; Aker, 2010). Recent papers study how intermediaries and market structure affect incidence and margins in agricultural supply chains (e.g., Bergquist and Dinerstein, 2020; Casaburi and Reed, 2022). We bring a climate-shock lens to these questions, documenting delayed but substantial downstream pass-through and a temporary widening of retailer margins, then relating margin compression to the timing of government ceilings. On the policy side, our evidence complements work on the (limited) effectiveness and unintended consequences of price controls in perishable-food markets (see classic theory in Stigler, 1952); we provide event-time evidence on when ceilings become binding and how compliance evolves during crisis.

Finally, we contribute a scalable empirical strategy for studying climate-induced supply shocks in data-scarce settings. In the absence of proprietary transaction or tracking data, we combine administrative market arrival records with remote sensing and geospatial tools to construct dynamic supply networks and measure shock exposure. Using satellite flood maps and georeferenced source locations, we infer source-destination routes via open-source road network algorithms. We also generate spatially explicit flood risk indices to bolster identification in our staggered event study framework. These tools offer a replicable approach for researchers studying climate vulnerability and market disruptions in low-data environments.

The remainder of this paper is organized as follows. Section 2 describes the study context. Section 3 presents the data sources. Section 4 outlines the empirical strategy. Section 5 presents the main findings. Section **??** shows the robustness of our results. Section 6 concludes.

2 Study context

2.1 Floods in Pakistan

Flooding in Pakistan can be broadly categorized into two types: fluvial and pluvial floods, with distinct spatial manifestations. Fluvial floods occur when rivers exceed their capacity to carry water and inundate adjacent floodplains. Pluvial flooding, by contrast, results from intense rainfall that exceeds local infiltration or drainage capacity; in Pakistan, this includes flash floods generated by hill torrents in foothill regions, as well as urban flooding in cities where drainage systems are inadequate.

The country has experienced three major floods in the last two decades—in 2011, 2014, and most recently, in 2022. The 2022 floods were driven by heavier-than-normal monsoon rains and unexpected torrential downpours, causing widespread inundation and river swelling. The impact was severe: critical infrastructure was damaged, livelihoods were disrupted, and nearly 6.4 million people were affected (United Nations, 2022a). The estimated economic losses amounted to $15.2 billion, or 4.31% of Pakistan's 2022 GDP. The agricultural sector in Sindh and Balochistan bore a disproportionate burden, with approximately 57% and 42% of agricultural lands in those provinces

affected, respectively (United Nations, 2022b,a).[1] Beyond immediate damage to standing crops and infrastructure, the floods may also have reduced investment in the subsequent planting cycle (Qamer et al., 2022).

To illustrate the extent of inundation, Figure 1 presents monthly satellite-based flood intensity maps from March to December 2022. Each panel represents conditions on the 30th of the month, using floodwater detection algorithms to assign intensity levels to 375-square-meter pixels. Blue areas represent low inundation, while red indicates complete flooding.[2]

The maps show early signs of flooding in April and May, with dramatic escalation in July and August—consistent with peak monsoon activity. Although the extent of flooding recedes from September onward, inundation levels remain above pre-monsoon levels by December, highlighting the persistence of flood impacts.

2.2 Agriculture Supply Chains in Pakistan

Agricultural commodities in Pakistan reach end consumers through multi-tiered supply chains. These can be divided into two main stages. In the first stage, produce moves from farms to government-regulated wholesale markets (*mandis*), where large lots are sold to wholesalers via public outcry auctions.[3] Farmers either sell to informal intermediaries who transport goods to *mandis*, or deliver produce directly. A few exceptions exist—for example, corporate buyers in the potato market or government grain procurement in wheat and rice. In the second stage, wholesalers and retailers disaggregate auctioned lots, with retailers transporting smaller quantities to local shops or street markets.

Our study focuses on the first stage: how floods in "source region" (farming areas) disrupt shipments to "destination districts" (consumer markets). We refer to source "regions" rather than districts because shipment data are recorded at sub-district level for origins and district level for destinations. We retain key food commodities, which are divided into three categories: vegetables (potatoes, onions, and tomatoes), fruits (apples and bananas), and grains & pulses (wheat, rice,

[1] An estimated 57% and 42% of agricultural lands in Sindh and Balochistan were affected by flood water, respectively.
[2] Further description of the flood data construction is provided in Section 3.5 and Appendix B.
[3] In a public outcry auction, the auctioneer announces opening, current, and winning bids. Auctions are open to all, and details are recorded by government officials.

Figure 1: Flood intensity March - December 2022

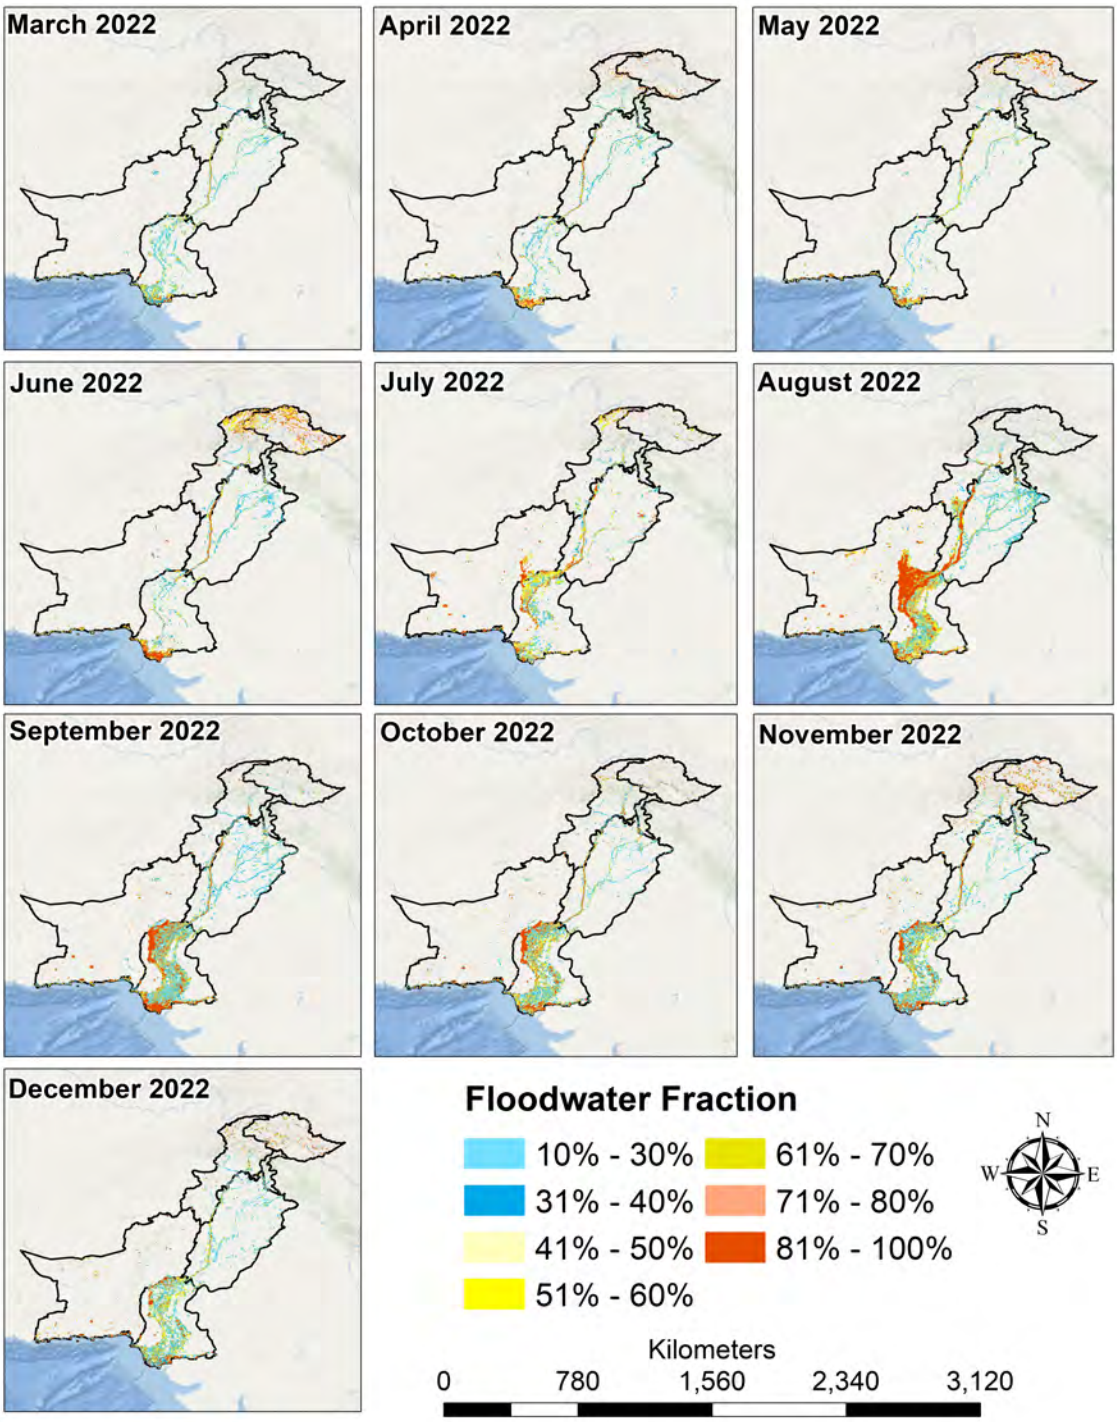

Note: Each map corresponds to the flood intensity recorded on the 30^{th} of each month from March to December of 2022. 'Percentage flooded' indicates the amount of floodwater in 375 square meters. The flood detection algorithm that generates these maps effectively removes cloud cover from flood maps, providing a clear view of flooded and dry areas. Intensities are authors' own calculation using source data from National Oceanic and Atmospheric Administration. Further details are provided in Appendix B.

maize, pulses). Additional details related to data and selection of commodities are provided in Section 3.

3 Data

This section describes the datasets underpinning our analysis of flood-induced disruptions to agricultural supply chains and market outcomes in Punjab. Drawing on administrative, geospatial, and price monitoring records from government and international sources, we compile a high-frequency panel covering crop flows, prices, regulatory ceilings, and flood exposure. These sources provide insights into supply volumes, wholesale and retail price dynamics, physical disruptions, and institutional responses. We also use district-level population figures from the 2020 Population Census compiled by the Pakistan Bureau of Statistics[4] to construct population weights in our destination-level regressions.

In what follows, we outline the six primary datasets used in our analysis. We organize the discussion by data function—tracking quantities, market prices, flood shocks, and enforcement constraints—and summarize key coverage and construction steps below.

3.1 Wholesale arrivals and supply volumes

We utilize high-frequency administrative data on agricultural commodity arrivals from the Department of Agriculture, Punjab. This dataset captures daily inflows of major crops into 36 regulated wholesale markets (destination districts) across Punjab, sourced from 223 sub-district-level regions nationwide (Figure 2). The commodities include vegetables (potatoes, onions, tomatoes), fruits (bananas, apples), and staple grains and pulses (wheat, rice, maize, and pulses).[5] We aggregate these daily records to the weekly level by source-destination-crop combinations for the period January to December 2022 to align with our flood data coverage.

To construct a balanced panel, we impute zeros for weeks with no recorded arrivals, yield-

[4] Available at: https://www.pbs.gov.pk/digital-census/detailed-results

[5] The crop classifications in our data are broad and do not distinguish between sub-types or grades within a commodity. For instance, while we observe data for apples, we cannot differentiate between varieties such as "Gala" and "Honeycrisp."

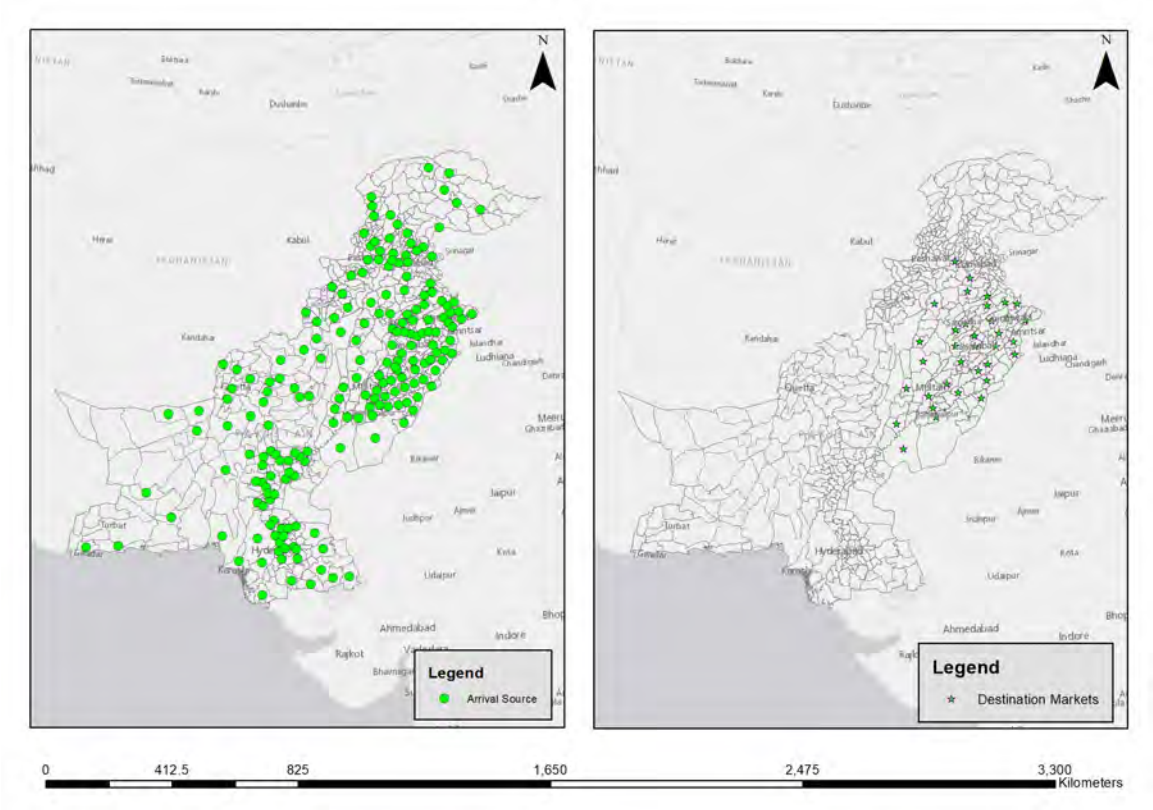

Figure 2: Locations of Source Regions and Destination Districts

Note: This figure shows the locations of 223 source regions and 36 destination districts spanning the study period.

ing over 218,000 crop-source-destination-week observations. These include 23,348 instances of strictly positive shipments. This high-frequency, spatially disaggregated dataset captures variation in supply volumes across crops, regions, and time, allowing us to track disruptions in commodity flows from producing areas to consumption centers before, during, and after the 2022 floods. Table 1 summarizes the dataset's coverage across crops, reporting the frequency of shipments and the average weekly supply per source based on weeks with non-zero arrivals.

While the dataset provides a detailed view of commodity flows into Punjab's formal wholesale markets, it does have certain limitations. It does not capture "out-of-wholesale-market" transactions, such as state-level grain procurement for food security and direct purchases by large grain mills. These are particularly relevant for staple grains like wheat and rice. Additionally, our data cover arrivals into Punjab only, and thus do not allow us to analyze supply dynamics or disruptions in the other provinces.

Table 1: Data Description

	Shipments	Supply	Auction price	Retail price	Price ceiling
Crop					
Apple	2,005	1,461.9	178.4	.	.
		(3,597.4)	(35.5)		
Banana	2,452	1,823.6	63.1	.	.
		(3,478.6)	(13.7)		
Citrus Fruit	980	2,188.3	72.7	.	.
		(4,630.7)	(90.9)		
Dates	372	116.8	136.4	.	.
		(287.1)	(20.6)		
Gram	932	1,347	181.3	.	.
		(2,180.4)	(34.4)		
Maize	251	7,095.2	82.6	.	.
		(10,684.7)	(73.5)		
Mango	563	2,262.6	128.9	.	.
		(5,369.5)	(26.2)		
Mash	798	820.9	254.1	.	.
		(2,256.6)	(63.5)		
Masoor	802	785	220.1	.	.
		(2,085.6)	(67.7)		
Moong	915	1,256.5	148	.	.
		(2,329.4)	(32.8)		
Onion	3,844	2,211	73.9	94.2	82.1
		(4,126.5)	(45.4)	(52.4)	(71.4)
Paddy	132	6,383.7	87	.	.
		(11,564.2)	(17.5)		
Potato	4,404	3,250.3	42.5	57.9	50.1
		(12,330.6)	(17.3)	(19.1)	(21.8)
Rice	1,066	2,470.3	170	.	.
		(3,701.5)	(45.7)		
Tomato	3,578	1,268.6	86.9	114.7	92.2
		(2,786.4)	(42.3)	(52)	(42.9)
Wheat	254	917.7	68.4	.	.
		(3,024)	(16.2)		
Total	23,348	2,030	103	85	71.5
		(6,409.6)	(71.6)	(48.2)	(51)

Note: The table reports the summary statistics for supply chain and prices data from January 2022 to December 2022 by each crop. 'Shipments' indicate the total weeks with at least one shipment from a source region during the study period. 'Supply' presents the average weekly supply per source in kilograms. Auction price, retail price, and the official price ceiling are all in PKR per kilogram. Values in parentheses are standard deviations. A dot (.) indicates no available data.

3.2 Wholesale auction prices

We supplement the supply chain data with auction price records from the same regulated wholesale markets across Punjab. These data from the Department of Agriculture, report the daily average price at which each crop is sold through an open outcry auction system. In this system, commission agents facilitate the auction process by inviting bids from traders, beoparis (local intermediaries), or, in some cases, farmers' representatives. The resulting prices, recorded in Pakistani Rupees (PKR) per kilogram, reflect prevailing wholesale market conditions at the point of first formal sale.

To ensure consistency with the supply data, we aggregate daily auction prices to the weekly level by crop and destination district. The dataset spans January to December 2022 and includes most crops in our sample, except in weeks or markets with no reported arrivals. In a few cases where prices were missing due to apparent data entry issues, we impute values using the average price for that crop in neighboring markets during the same week.

The 'Auction price' column in Table 1 summarizes the distribution of these prices across crops, based on weeks with non-zero shipments. These prices serve as a key outcome in our analysis, enabling us to study how flood-induced supply disruptions shaped price dynamics across time and space.

3.3 Retail prices

To capture consumer-level price dynamics, we use retail price data collected by the Pakistan Bureau of Statistics (PBS) as part of the government's price monitoring and control efforts. These data report weekly prices (in PKR per kilogram) for key perishable commodities from eight major urban centers in Punjab: Bahawalpur, Faisalabad, Gujranwala, Lahore, Multan, Rawalpindi, Sargodha, and Sialkot.[6] Prices are collected through physical visits by enumerators and reflect prevailing consumer-facing rates in local markets.

The retail price data are available only for three major vegetables, onions, potatoes, and tomatoes, which are core components of the food basket and highly sensitive to short-term supply

[6] The system is called the Decision Support System for Inflation (DSSI). It provides data for essential food items across 17 major cities in the country and includes prices of fresh vegetables (potatoes, tomatoes, and onions). It does not provide data for fruits or grains.

shocks. The dataset includes 43,656 crop-district-week observations across the eight cities, with coverage split roughly evenly across the three crops.

These data are not publicly accessible and were obtained through personal contacts within the government. The Retail price column in Table 1 summarizes the average reported prices across the three covered crops. Retail prices help us assess how wholesale disruptions during the 2022 floods affected final consumer prices in urban areas.

3.4 Government price controls

To analyze the role of state intervention in consumer price stabilization during the flood period, we compile weekly price ceiling data, referred to as Deputy Commissioner (DC) rates, for three key vegetables: onions, potatoes, and tomatoes. These ceilings represent official maximum retail prices set by Market Committees under the oversight of the district administration. The system is intended to protect consumers from price spikes and inflationary pressures, especially for essential perishable commodities. In practice, however, price ceilings often diverge from real-time market conditions and are implemented with varying effectiveness across districts.

We obtain weekly DC rates for the same eight urban districts covered by the retail price data (Bahawalpur, Faisalabad, Gujranwala, Lahore, Multan, Rawalpindi, Sargodha, and Sialkot) and compiled by PBS. We match the weekly ceiling prices to corresponding retail prices by crop, district, and week to examine whether ceilings became binding and to assess their potential impact on market margins.

Price ceilings face weak enforcement and rent-seeking (Naz and Naz, 2024). At the wholesale (mandi) level, large traders often bypass DC rates, citing their misalignment with actual supply-demand conditions. Retailers similarly report limited adherence, pointing to higher operational costs and lax enforcement. In contrast, those who comply adjust product quality or reduce quantities to offset margin losses.

The Price ceiling column in Table 1 summarizes the variation in official ceilings across the three crops. We use this information in our analysis to study the interaction between public price regulation and market behavior during the flood period, specifically, whether ceilings constrained

retail price adjustments and how margins evolved in response to supply shocks.

3.5 Observed flood exposure

To capture flood exposure across source regions during 2022, we rely on satellite-based inundation data from the National Oceanic and Atmospheric Administration (NOAA), specifically from the Suomi-NPP and NOAA-20 VIIRS (Visible Infrared Imaging Radiometer Suite) satellites. These data provide near-daily coverage with a spatial resolution of 375 meters, allowing for fine-grained tracking of surface water anomalies at the sub-district level.

We obtain flood maps from January 01 to December 31, 2022. After downloading the raw pixel-level raster files, we preprocess them to match the administrative boundaries of source regions in our supply chain dataset. Each pixel is coded as an 8-bit integer (values from 0 to 255), representing surface conditions.[7] To minimize false positives (e.g., temporary waterlogging from rainfall), we define flood presence as continuous water detection over at least 10 consecutive days. For each week, we compute the proportion of flooded area within a source region and transform this to a standardized flood intensity measure, making values comparable across space and time.[8] We also classify a source region as flooded if more than 10 percent of the source area was flooded during the week.[9]

Figure 3 illustrates pixel-level flood intensity for the representative source region of Badin, based on satellite imagery. Each panel shows monthly floodwater fractions across the district from March to December 2022. Flood intensity is measured as the proportion of a 375m pixel covered by surface water and is color-coded in bands ranging from 10–30% (light blue) to 81–100% (dark red). The maps highlight both the temporal evolution and spatial extent of inundation within the region. Notably, significant flooding becomes visible in July and peaks in August–September, particularly in the southern and coastal parts of the district. These visualizations help validate the construction of weekly flood intensity measures used in the analysis.

[7] Appendix B details the pixel-to-region conversion process.

[8] Flood intensity is calculated as: $\frac{FloodedArea_{st}}{TotalArea_s}$

[9] We adopt a 10 percent threshold to balance sensitivity and specificity in identifying meaningful flood exposure. Smaller thresholds (e.g., 1–5%) may overstate flood events, while much higher thresholds risk missing significant but spatially concentrated inundation that could still disrupt agricultural operations. The 10 percent benchmark reflects a conservative but practical cutoff for operational relevance.

Figure 3: Flood intensity in District Badin from March - December 2022

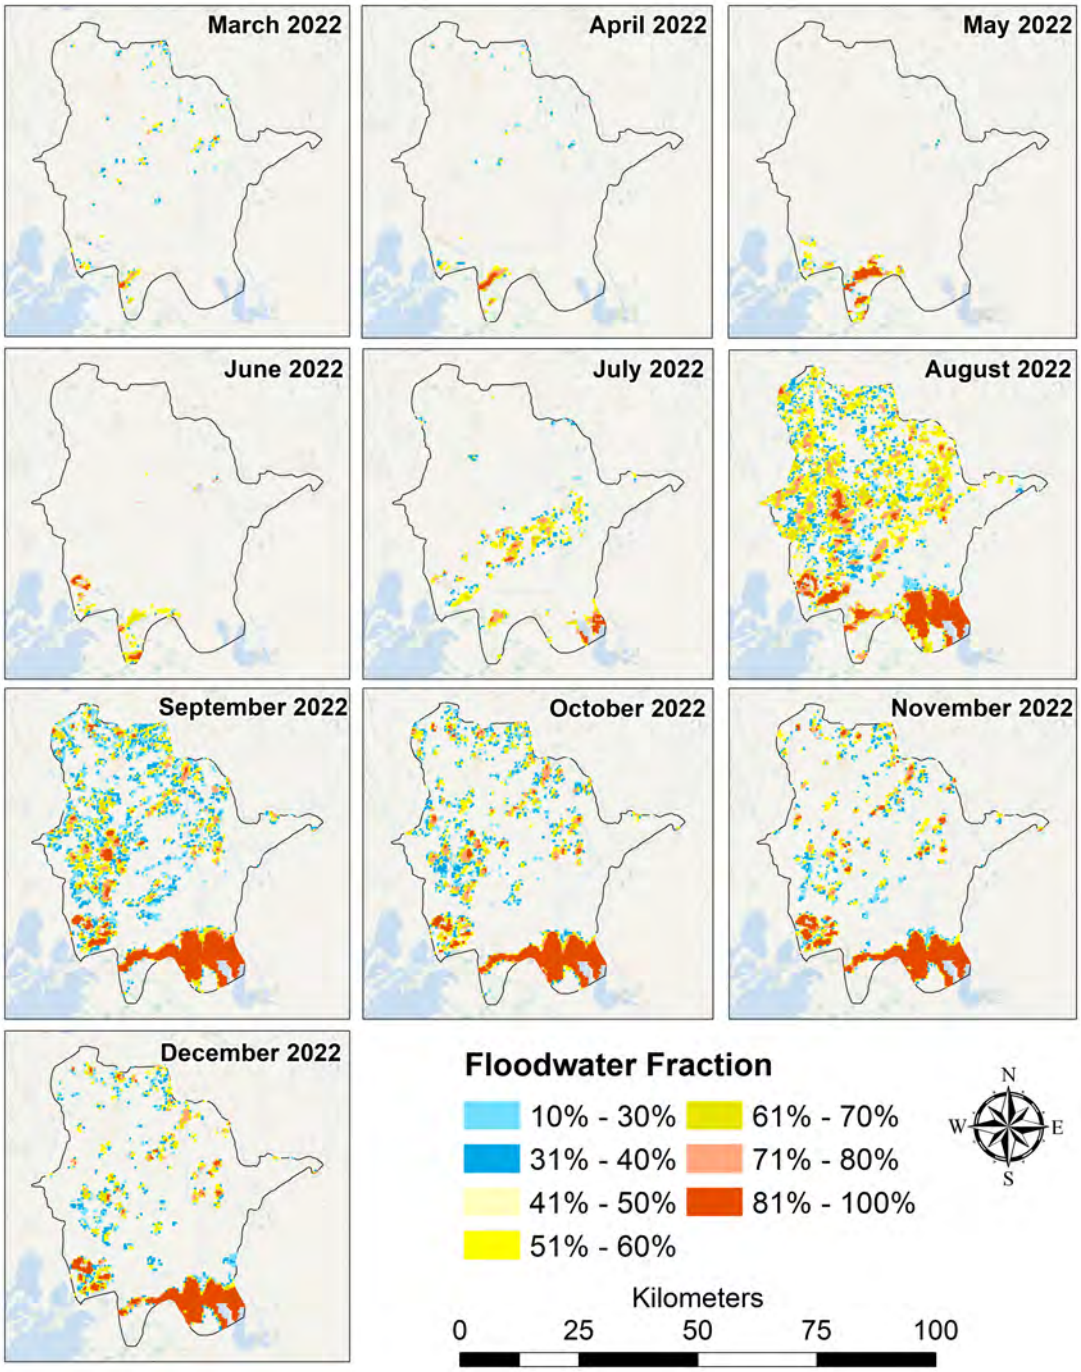

Note: Each map corresponds to the flood intensity recorded on the 30^{th} of each month from March to December 2022. 'Percentage flooded' indicates the amount of floodwater in an area of 375 square meters. The flood detection algorithm used in generating these maps removes cloud cover from flood maps, providing a clear view of flooded and dry areas. Intensities are authors' own calculation using source data from NOAA

Figure 4, Panel (a) shows that around 90% of source regions recorded their first flood between July and August (weeks 26–36), consistent with Pakistan's monsoon season. The distribution of maximum standardized flood intensity across regions is shown in Panel (b) of Figure 4, with the sharpest surge observed from late August to September (weeks 33–39), coinciding with widespread media coverage of flood devastation.

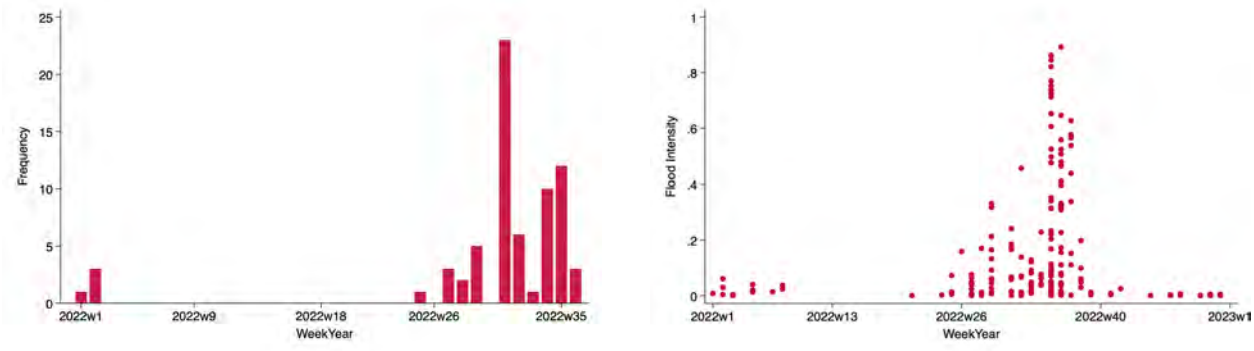

(a) First recorded instance of flooding **(b)** Week of maximum flood intensity

Figure 4: Temporal patterns of flooding across source regions in 2022

Notes: Panel (a) shows the number of source regions by the week of their first recorded flood detection. Panel (b) shows the week in which each source region experienced its highest standardized flood intensity, with each dot representing a source region. A source is classified as flooded in a given week if more than 10% of its area is covered by floodwater based on satellite-detected inundation.

3.6 Underlying flood risk

In addition to observed flood exposure during 2022, we construct a spatial flood risk index to capture the underlying susceptibility of each region to flooding. This static measure of risk, derived from geospatial and historical characteristics, serves to explore heterogeneity in anticipatory responses across low- and high-risk areas. While our main results are based on actual inundation, this pre-existing risk profile helps distinguish between regions with historically higher vulnerability and those where flooding was less expected.

We generate the flood risk profile using a multi-criteria approach implemented through Google Earth Engine (GEE). The methodology integrates several physical and hydrological parameters: (i) historical water occurrence from the Global Surface Water dataset (Pekel et al., 2016), (ii) proximity to permanent water bodies, (iii) topography based on Shuttle Radar Topography Mission

(SRTM) Digital Elevation Model data, (iv) long-term precipitation trends, (v) terrain morphology via the Topographic Position Index (TPI), and (vi) vegetation density and wetness derived from remote sensing imagery. These factors are commonly used in the flood risk assessment literature and enable spatially explicit mapping at high resolution (e.g., Swain et al., 2020; Moazzam et al., 2018).

Each pixel in the national grid is assigned a flood hazard score ranging from 1 (very low risk) to 5 (very high risk). For each source region, we calculate the average hazard score by aggregating the pixel-level values within its boundaries. The resulting region-level scores provide a continuous measure of baseline flood risk. Additional technical details on data sources and processing steps are provided in Appendix C. Figure 5 displays the spatial distribution of calculated flood risk across Pakistan.

While our flood risk index aligns with historical flood-prone zones, such as areas near the Indus River or around major lakes, some regions with high risk did not experience flooding in 2022, and vice versa. This divergence reflects the inherent difference between probabilistic risk estimates and realized flood events. We leverage this variation in our heterogeneity analysis to better understand the role of perceived vulnerability in shaping behavioral responses.

3.7 Supply road networks

We construct a geospatial dataset of optimal road routes linking source regions to destination markets. For each of the 8028 source-destination pairs (combinations of 223 unique source regions and 36 destination districts), we calculate the optimal travel path based on driving time using the RoutingPy Python library and the Mapbox Navigation Services API (Viktorovic, 2023; Rzeszewski, 2023). Among the four available travel modes, we select "driving with traffic" to obtain realistic travel time estimates and select the route with the shortest travel duration as the optimal path. Figure 6 visualizes these computed routes across Punjab.

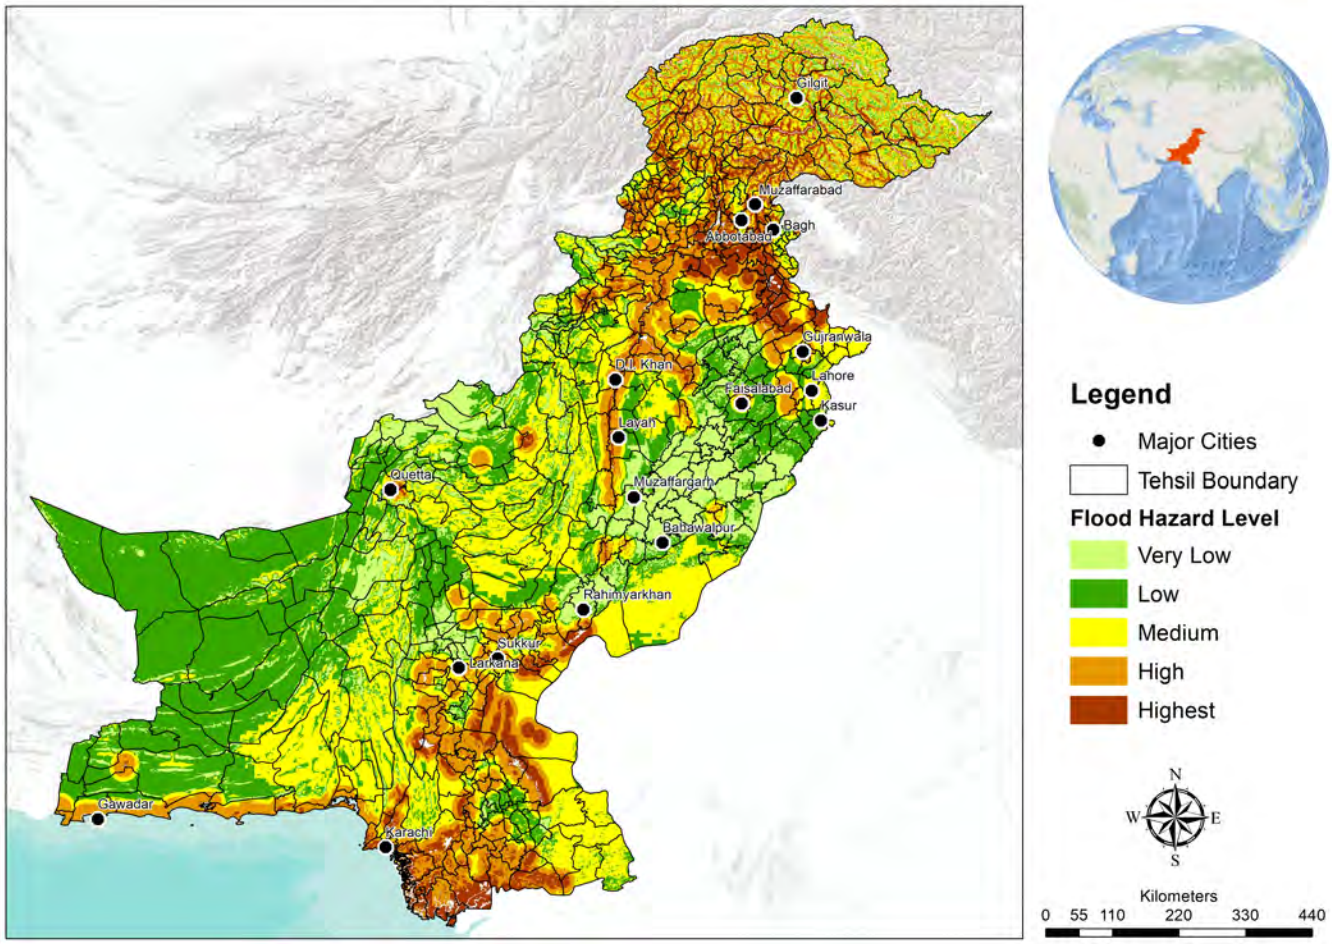

Figure 5: Spatial flood risk for all regions in Pakistan

Notes: This map displays spatial flood hazard levels for all sub-districts (tehsils) in Pakistan, based on a composite index constructed using elevation, proximity to water bodies, historical water occurrence, precipitation trends, terrain features, and vegetation indicators. Flood risk is classified into five levels—Very Low, Low, Medium, High, and Highest—represented by increasing color intensity from light green to dark brown. Risk scores are averaged over all pixels within each tehsil boundary. Major cities are marked for reference. The map illustrates both riverine flood-prone zones (e.g., along the Indus River) and locally elevated risks near lakes or coastal areas.

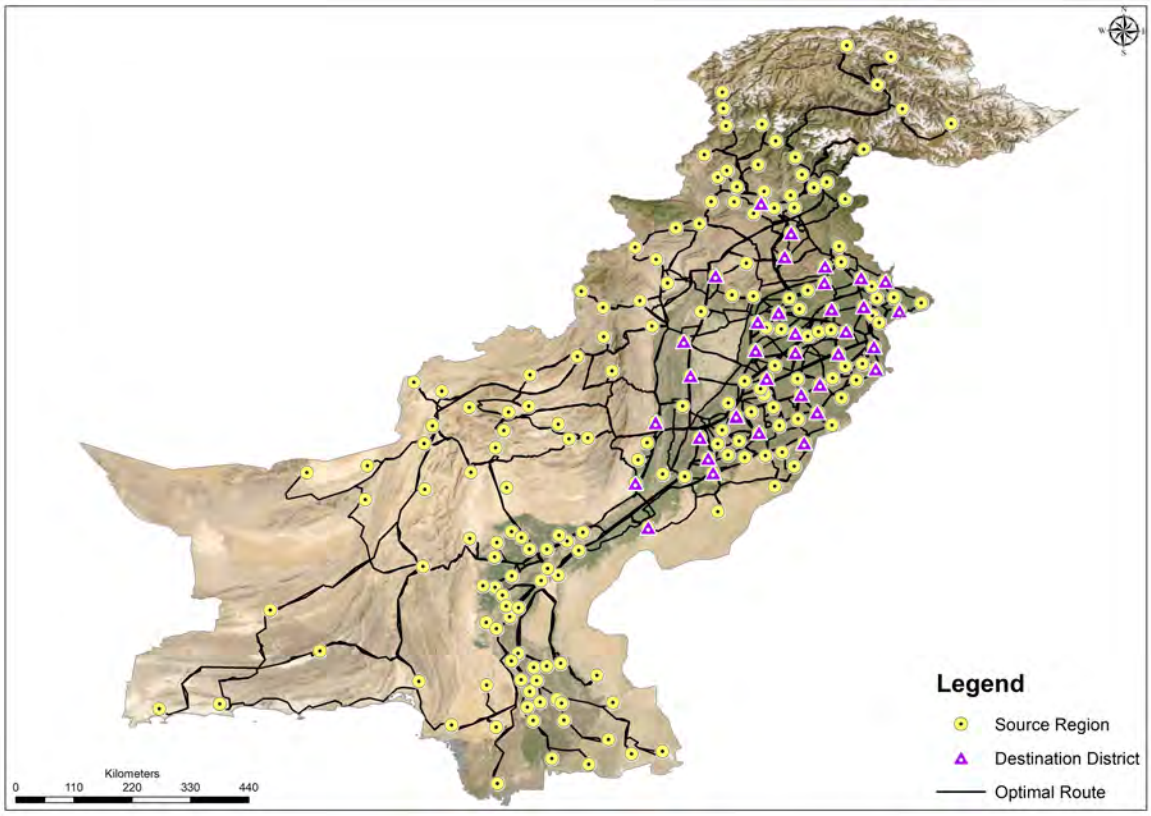

Figure 6: Optimal routes connecting Source Regions with Destination Districts

To measure the impact of flooding on transport infrastructure, we spatially overlay daily flood maps onto Pakistan's road network. For each source-destination pair, we generate a 375-meter buffer around the optimal route—matching the spatial resolution of the flood data—and identify any floodwater detected within that buffer. This allows us to detect when and where a road segment intersects with flooded terrain. Based on this overlay, we construct two key variables: (i) a binary indicator that equals 1 on days when any part of the route was inundated, and (ii) a continuous route-level flood intensity measure, calculated as the proportion of the buffered area that overlaps with flood pixels.

Figure 7 illustrates this overlay using flood data from August 28, 2022. The map shows both the full road network (black lines) and areas under floodwater (red pixels). The inset highlights a specific route from Kachhi (source) to Rahim Yar Khan (destination), which shows no direct flooding at either endpoint, but significant inundation along the connecting road, suggesting likely disruption from route-level flooding even in the absence of source or destination flooding.

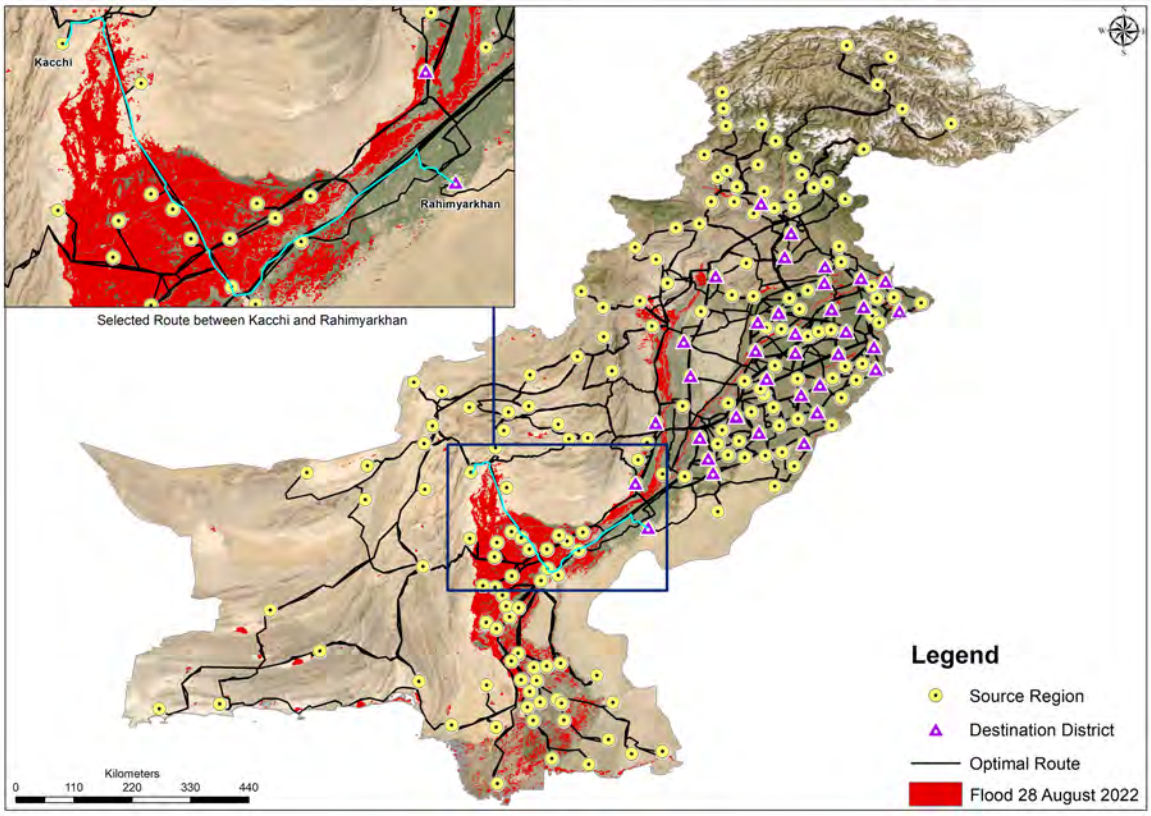

Figure 7: Impact of August 28, 2022 flooding on a sample optimal road network route.

Notes: This map overlays flood extent on August 28, 2022 (in red), with optimal source-to-destination road routes (black lines) computed using Mapbox navigation data. Yellow circles represent agricultural source regions, and purple triangles denote destination districts. The inset highlights a specific route between Kachhi and Rahim Yar Khan (in cyan), showing that although neither endpoint was flooded, the connecting route experienced substantial inundation.

These data have three limitations. First, the road network is based on static data from 2021, and we cannot account for structural damage to road segments caused by floods. If a road were physically destroyed mid-year, we would not observe this change. Second, we do not observe real-time routing decisions by transporters—they may have rerouted around flooded areas or experienced delays, which we cannot capture. Finally, our analysis focuses on inter-district roads and does not cover small intra-district or rural roadways, limiting our ability to fully disentangle the effects of source-region flooding from those of route disruption when both occur in the same administrative district.

4 Empirical Methodology

Our empirical strategy proceeds in two steps. First, we estimate source-level impacts of flooding on crop supplies using a doubly-robust, staggered difference-in-differences event-study estimator (Sant'Anna and Zhao, 2020). Second, we measure how these shocks propagate to downstream markets via a shift–share exposure instrument and the same event-study framework.

4.1 Source-Level Impacts

To quantify how flooding disrupts agricultural supply at the origin, we estimate a dynamic difference-in-differences event-study model using the doubly-robust estimator of Sant'Anna and Zhao (2020). We define y_{cst} as the total quantity of crop c shipped from source region s in calendar week t. Treatment onset t_s^{flood} is defined as the first week in which satellite-derived flood intensity exceeds our threshold in region s. Our baseline specification is

$$\ln(y_{cst}) = \alpha_{cs} \; + \; \lambda_t \; + \; \sum_{k=-K}^{\hat{K}} \beta_k \, \mathbf{1}\{\, t - t_s^{\text{flood}} = k\} \; + \; \varepsilon_{cst}\,, \tag{1}$$

where α_{cs} are crop–source fixed effects, capturing all time-invariant heterogeneity (soil quality, market access, typical cropping patterns), and λ_t are week–year fixed effects, absorbing common seasonal and macro-level shocks. The event-study coefficients β_k trace the percentage change in supply k weeks before or after the initial flood week, with the omitted category $k = -1$ normalizing to the week immediately prior to inundation. To balance precision and dynamic flexibility, we report estimates for $K = 10$ and $\hat{K} = 20$. The term ε_{cst} is an idiosyncratic error that captures remaining, time-varying shocks to supply – such as local demand shifts or unobserved logistical constraints – not accounted for by the fixed effects or event-time indicators.

Estimation is weighted by the geographic area of each source region to reflect policy relevance at that level, and standard errors are clustered at the source region to allow for arbitrary serial correlation. Identification relies on the assumption that, in the absence of flooding, ever-flooded and never- (or not-yet-)flooded source regions would have followed parallel trends in $\ln(y_{cst})$, an assumption we verify by inspecting pre-treatment estimates of β_k for statistical insignificance.

In robustness checks, we also interact each event-time indicator with the continuous flood-

intensity measure g_{st} to confirm that larger shocks produce proportionally greater supply declines.

4.2 Destination-Level Impacts

To trace how upstream flooding propagates to consumer markets, we construct a shift–share exposure measure for each crop–destination pair, à la Borusyak et al. (2025), and estimate a dynamic event-study model in logs. Let y_{cdt} be the outcome (quantity or price) of crop c at destination market d in calendar week t. We first compute an instrument

$$z_{cdt} \;=\; \sum_{s=1}^{S} s_{csdw}\; g_{st},$$

where s_{csdw} is the historical, pre-flood share of crop c that market d sources from region s in calendar week $w = \mathrm{WeekOfYear}(t)$, averaged over baseline years 2018–2022, and g_{st} is the flood intensity ($\in [0,1]$) in region s at week t. We define the initial exposure date as

$$t^{\exp}_{cd} \;=\; \min\{\, t : z_{cdt} > 0\},$$

the first week in which market d experiences any upstream flooding of crop c.

Using this exposure date, we then estimate

$$\ln(y_{cdt}) \;=\; \alpha_{cd} \;+\; \lambda_t \;+\; \sum_{k=-K'}^{K''} \gamma_k\, \mathbf{1}\{\, t - t^{\exp}_{cd} = k\} \;+\; \nu_{cdt}, \tag{2}$$

where α_{cd} are crop–destination fixed effects capturing time-invariant differences in market size and preferences, λ_t are week–year fixed effects absorbing common seasonality and aggregate shocks, and the coefficients γ_k trace the percent deviation in arrivals k weeks relative to the first exposure (omitting $k = -1$ as the reference period). To balance flexibility and precision, we report estimates for $K' = 10$ and $K'' = 20$.

Observations are weighted by the population of destination d, reflecting our interest in welfare-relevant market impacts, and standard errors are clustered at the destination level to allow for arbitrary serial correlation in market-level shocks. Identification requires that, absent flooding, exposed and never-exposed markets would have followed parallel trends in $\ln(y_{cdt})$; we assess this by verifying that the estimated γ_k for $k < 0$ are statistically indistinguishable from zero.

In robustness checks, we re-estimate equation (2) on the subset of ever-exposed markets; vary the definitions of the event window; and interact the event-time indicators with the continuous exposure measure z_{cdt} to confirm that larger upstream shocks yield proportionally greater downstream disruptions. As we show in section **??**, these alternative specifications produce consistent results, underscoring the robustness of our findings.

4.3 Transport–Network Channel: Road Inundation

Beyond damaging production at origins, floods can sever the transport links that connect sources to downstream markets. To incorporate this logistical channel, we exploit the route panel described in Section 3.7. For each source–destination pair (s, d) we precompute the optimal ordinary-time path on the road network and overlay daily flood masks within a 375 m corridor around that path. We then aggregate these daily overlays to the week–year index t, used throughout the paper, by taking the within-week maximum inundation.[10] This produces two route-level measures: an indicator I^{route}_{sdt} equal to one if any part of the buffered route is inundated on any day of week t, and a continuous flooded-share $F_{sdt} \in [0, 1]$ equal to the proportion of the buffered area intersecting flood pixels during week t.

We first estimate a route-level event study that isolates connectivity shocks from contemporaneous shocks at origins and destinations. Let y_{csdt} denote shipments of crop c from s to d in week t. Define the first exposure date for the route as $E_{sd} = \min\{t : I^{\text{route}}_{sdt} = 1\}$. We estimate

$$\ln(y_{csdt}) \;=\; \alpha_{csd} \;+\; \alpha_{st} \;+\; \alpha_{dt} \;+\; \sum_{\tau=-L}^{U} \theta_\tau \, \mathbf{1}\{t - E_{sd} = \tau\} \cdot \widetilde{F}_{sdt} \;+\; \xi_{csdt}, \tag{3}$$

where α_{csd} are route (crop–source–destination) fixed effects that absorb all time-invariant pair characteristics, while α_{st} and α_{dt} are source–week and destination–week fixed effects that absorb *all* time-varying shocks at the endpoints, including local flooding and demand shifts. The term $\widetilde{F}_{sdt}$ is either the binary I^{route}_{sdt} or the continuous F_{sdt} (scaled to $[0, 1]$). The coefficients θ_τ trace the dynamic deviation in link-specific shipments around the first route inundation, with $\tau = -1$ omitted. Identification therefore comes from within-route variation in inundation, holding fixed any shocks at s or d. We report specifications with $L = 10$ and $U = 20$ weeks; standard errors

[10] Results are robust to using weekly means rather than maxima.

are clustered at the (s, d) route level. The error term ξ_{csdt} captures idiosyncratic route shocks not absorbed by the fixed effects or event-time terms.

To allow the road network to mediate downstream exposure, we also make the destination event study network-aware by reweighting upstream shocks according to contemporaneous reachability. Let s_{csdw} be the historical pre-flood sourcing share of crop c from s to d in calendar week $w = \mathrm{WeekOfYear}(t)$ (averaged over 2018–2022). We define a simple accessibility factor from the flooded-share,

$$A_{sdt} \;=\; 1 - F_{sdt} \quad \in [0, 1],$$

so that A_{sdt} downweights sources whose ordinary-time path to d is inundated in week t. The network-adjusted weights are then

$$\tilde{s}_{csd,t} \;=\; \frac{s_{csdw}\, A_{sdt}}{\sum_u s_{cudw}\, A_{udt}},$$

and the corresponding exposure measure becomes

$$\tilde{z}_{cdt} \;=\; \sum_{s=1}^{S} \tilde{s}_{csd,t}\, g_{st}, \tag{4}$$

where $g_{st} \in [0, 1]$ is source-level flood intensity. We define the network-adjusted exposure onset as $t_{cd}^{\mathrm{exp,net}} = \min\{t : \tilde{z}_{cdt} > 0\}$ and re-estimate equation (2) replacing t_{cd}^{exp} with $t_{cd}^{\mathrm{exp,net}}$. In a complementary dose–response specification we also include $\tilde{z}_{cdt}$ directly:

$$\ln(y_{cdt}) \;=\; \alpha_{cd} + \lambda_t \;+\; \eta\, \tilde{z}_{cdt} \;+\; \sum_{\tau \neq -1} \gamma_\tau\, \mathbf{1}\{t - t_{cd}^{\mathrm{exp,net}} = \tau\} \;+\; \nu_{cdt}, \tag{5}$$

where η summarizes contemporaneous pass-through of network-adjusted upstream shocks into market outcomes and ν_{cdt} captures remaining idiosyncratic market-level shocks. As in Section 4.2, observations are weighted by destination population and standard errors are clustered at the destination level.

Finally, to gauge how much of the source-level supply loss reflects impaired egress rather than crop damage, we augment the source event study in equation (1) with an origin-specific measure of network disruption. We construct R_{st} as the egress-weighted average flooded-share across all

routes emanating from s in week t,

$$R_{st} \;=\; \sum_{d} \omega_{sdw}\, F_{sdt}, \qquad \omega_{sdw} \;=\; \frac{\sum_{c} s_{csdw}}{\sum_{d'} \sum_{c} s_{csd'w}},$$

and estimate

$$\ln(y_{cst}) \;=\; \alpha_{cs} + \lambda_t + \sum_{k \neq -1} \beta_k\, \mathbf{1}\{t - t_s^{\text{flood}} = k\} + \sum_{k \neq -1} \delta_k\, \mathbf{1}\{t - t_s^{\text{flood}} = k\} \cdot R_{st} + \varepsilon_{cst}. \quad (6)$$

The sequence $\{\delta_k\}$ quantifies the incremental contribution of within-source connectivity constraints to supply losses, conditional on the source's flood intensity g_{st}. Stability of $\{\beta_k\}$ upon inclusion of R_{st} would indicate that production damage dominates; attenuation would imply a meaningful role for road closures.

Two caveats from the data construction inform interpretation. First, the underlying routable network is static (2021), so permanent flood-induced damage mid-year is not observed; our measures capture inundation rather than structural washouts. Second, we do not observe real-time re-routing choices. Using weekly maxima for F_{sdt} and I_{sdt}^{route} therefore delivers an upper bound on route exposure in a given week but preserves the quasi-experimental timing necessary for the event-study design. In Section **??**, we show that results are robust to alternative aggregation rules and to excluding pairs that experience repeated episodes of route inundation within a short horizon.

5 Results

This section presents the empirical estimates of how extreme flooding affected agricultural supply chains across Punjab. We organize the analysis in two parts. First, we examine the source side to assess how flood exposure altered the quantity of crops supplied from production regions to markets, using a dynamic event study approach as outlined in Section 4.1. Second, we explore the downstream effects on destination markets by estimating how exposure to disrupted supply sources shaped quantities received, price levels, retailer margins, and policy enforcement outcomes using the methodology outlined in Section 4.2. Throughout, we disaggregate results by crop type and leverage spatial variation in flood risk to examine heterogeneity in impacts.

5.1 Source-Level Impacts on Quantity Supplied

We begin by estimating how flooding in source regions affected the quantity of crops supplied to wholesale markets. Figure 8 displays dynamic treatment effects from our event study specification, where the outcome is the log quantity of crops supplied from a source region in a given week. The estimates show no significant deviation from baseline trends in the pre-treatment period, supporting the parallel trends assumption. A sharp and immediate decline in supply begins at the onset of flooding (week 0) and persists for at least 20 weeks. The average treatment effect peaks at a reduction of nearly 50 percent, suggesting substantial and lasting disruption to production or transport capacity in affected regions.

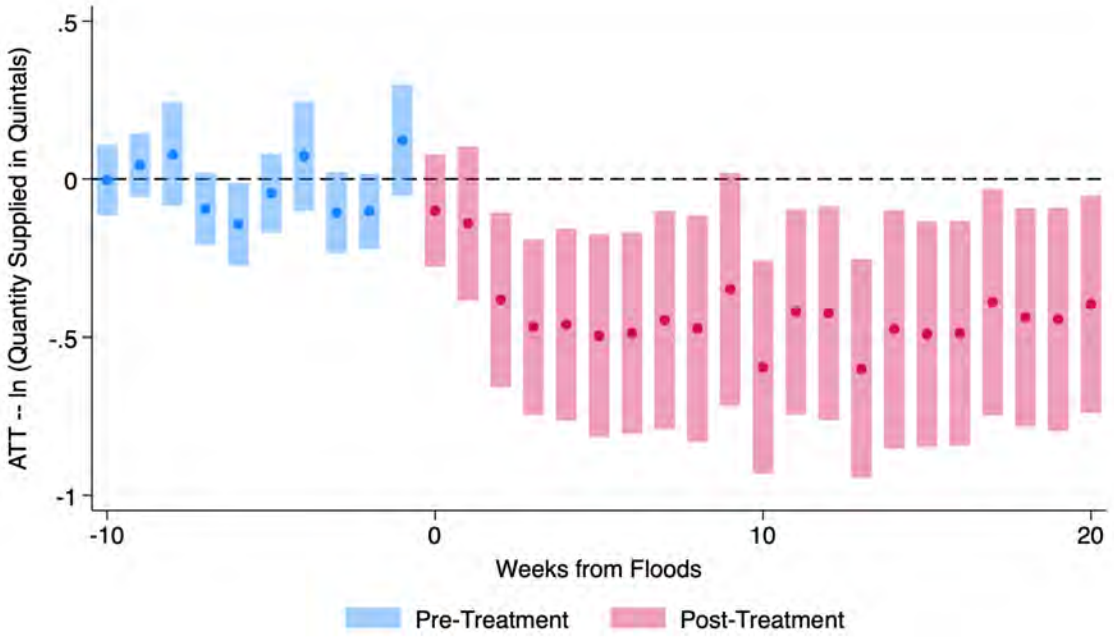

Figure 8: Flood Impact on Quantity Supplied from Source Regions

Notes: This figure plots the dynamic treatment effects of flooding on the log quantity of crops supplied from source regions using an event study specification. The vertical axis shows the average treatment effect on the treated (ATT) in log units of quintals supplied, and the horizontal axis denotes weeks relative to the first week of observed flooding in a source region (week 0). Estimates are weighted by source area and include crop–source and time (week-year) fixed effects. The pre-treatment period (weeks -10 to -1) is shaded in blue, and the post-treatment period (weeks 0 to 20) is shaded in pink. The plotted points represent point estimates, and the shaded bars show 95% confidence intervals. The results indicate a sharp decline in supply from flooded regions beginning in the week of flood onset, with persistent effects that extend at least 20 weeks post-event.

To investigate heterogeneity in supply disruptions, Figure 9 disaggregates the dynamic treatment effects by crop category.[11] Both fruits and vegetables exhibit a significant decline in quantities supplied following the onset of flooding. However, the timing and severity of the response differ: vegetables show an immediate and sharp drop beginning in the week of flood exposure, while the impact on fruits is slightly delayed and less steep. This distinction likely reflects differences in perishability and harvest logistics—vegetables such as potatoes and tomatoes are more sensitive to inundation and have shorter post-harvest lifespans, leaving farmers with little flexibility in the face of flooding. Fruits, particularly those harvested from trees, may have slightly more buffer time, moderating the immediacy of the disruption.

Grains, by contrast, show no meaningful deviation from baseline trends, consistent with their longer shelf life, lower perishability, and the possibility of temporary on-farm or centralized storage. In many cases, grain procurement and distribution are also mediated by government programs, which may insulate their market flow from short-run climatic shocks.[12] These differences highlight how the biophysical characteristics of commodities—especially perishability and storability—shape the transmission of extreme weather events through supply networks.

In Figure 10, we exploit our flood-risk index to compare supply responses in high-risk versus low-risk source regions. Across both high- and low-hazard areas, pre-treatment coefficients remain tightly centered around zero, confirming parallel trends and showing little evidence of anticipatory harvesting, even where historical flood risk is greatest. Once flooding begins (week 0), supply falls sharply in both groups. The depth and duration of that collapse differ markedly: low-hazard regions experience a larger initial drop—over 50 percent at its peak—and a slower recovery, whereas high-hazard regions see a more moderate decline (around 30–40 percent) and a faster rebound toward the baseline. This pattern suggests that prior exposure to flooding attenuates the marginal impact

[11] Further disaggregated impacts by crop are presented in Figure A3.

[12] In Pakistan, wheat procurement and distribution are coordinated through both provincial and federal agencies. Each province sets its own procurement targets and support prices based on local production conditions and fiscal space. In Punjab and Sindh, the Food Departments are primarily responsible for purchasing wheat from farmers at government-set prices, storing it in public warehouses, and releasing it to flour mills throughout the year to stabilize supply and prices. The federal Pakistan Agricultural Storage and Services Corporation (PASSCO) steps in when provinces are unable to meet their procurement targets or when strategic reserves require replenishment. For instance, after Sindh fell short of its 1.4 million tonne wheat procurement target in 2022, the federal government authorized PASSCO to procure additional stock in the province to avert shortfalls in public distribution. See: Dawn (2023), "*Agriculture: The muddled affairs of the food department*," July 24.

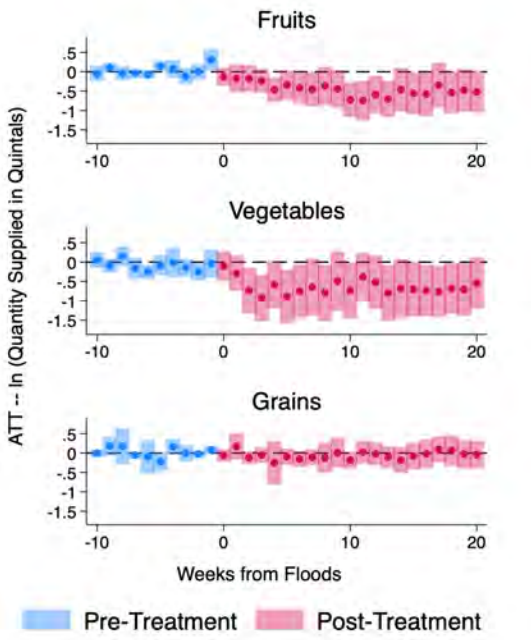

Figure 9: Flood Impact on Quantity Supplied, by Crop Type

Notes: This figure presents the dynamic treatment effects of flooding on the log quantity of crop supplied from source regions, disaggregated by crop category: fruits, vegetables, grains, and all commodities. Estimates are from a doubly robust event study model with crop–source and week fixed effects and are weighted by the geographic area of each source region. The horizontal axis represents weeks from the first flood exposure (week 0). Blue bars indicate the pre-treatment period; pink bars reflect the post-treatment period. Points represent coefficient estimates, and bars show 95% confidence intervals.

of subsequent shocks (a result consistent with Patel (2025)). These contrasts imply that farmers in high-risk zones may have developed adaptive strategies (such as more robust drainage, earlier planting cycles, or informal storage) that potentially mitigate the shock. In contrast, farmers in low-risk zones, less accustomed to such shocks, may face more severe and unanticipated disruptions.

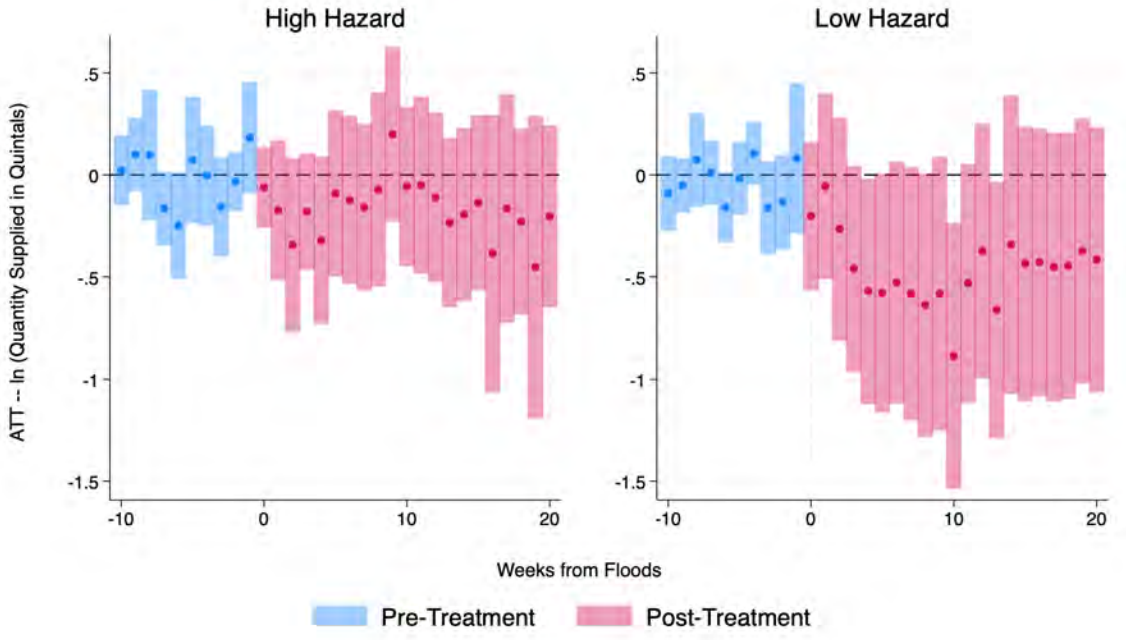

Figure 10: Flood Impact on Quantity Supplied, by Flood Hazard

Notes: This figure shows dynamic treatment effects of flooding on the log quantity supplied from source regions, stratified by ex-ante flood hazard. "High Hazard" (left) and "Low Hazard" (right) panels present event-study coefficients from the doubly robust DiD model (crop–source and week fixed effects; area weights). The horizontal axis measures weeks from the first observed flooding (week 0). Shaded bars denote 95% confidence intervals: blue for pre-treatment (weeks –10 to –1) and pink for post-treatment (weeks 0 to 20). Point estimates illustrate how supply responds over time in regions with different inherent flood risk.

In summary, the event-study estimates in Figures 8–10 reveal that flooding in source regions triggers an immediate and lasting collapse in supply—particularly for perishable vegetables—while fruits suffer a smaller, slightly lagged decline and grains remain largely unaffected. Regions with chronic flood risk appear somewhat more resilient, experiencing shallower drops and quicker recovery than low-hazard areas. Together, these patterns underscore how crop perishability and historical exposure shape upstream vulnerability.

Having documented the magnitude and heterogeneity of supply disruptions at the source, we now turn to the downstream consequences: how these shocks affect the quantities arriving in

destination markets, auction and retail prices, and intermediary margins.

5.2 Destination-Level Impacts on Supply and Prices

We now turn to how upstream supply shocks translate into market-level outcomes. In this subsection, we measure each destination's exposure to flooding by combining the share of its usual supply coming from flooded sources with the contemporaneous flood intensity, and then trace the dynamic effects of that exposure on (i) the quantity of each crop actually arriving in wholesale markets and (ii) both auction and retail prices. We ask: when source regions flood, how much less crop reaches spot markets, how quickly do prices adjust at the auction and retail stages, and to what extent do intermediary margins and price-ceiling enforcement respond? These estimates will show to what degree downstream markets absorb, amplify, or mitigate the upstream disruption.

5.2.1 Quantity Supplied

We next examine how source-level flooding translates into actual deliveries to the wholesale markets. Figure 11 shows the ATT across all commodities: there is no significant departure from baseline in the first two post-flood weeks, confirming that supply holds steady immediately after upstream inundation. However, beginning in week 3, delivered volumes start to fall, reaching a trough of approximately –1.0 log points (65%) by week 15.

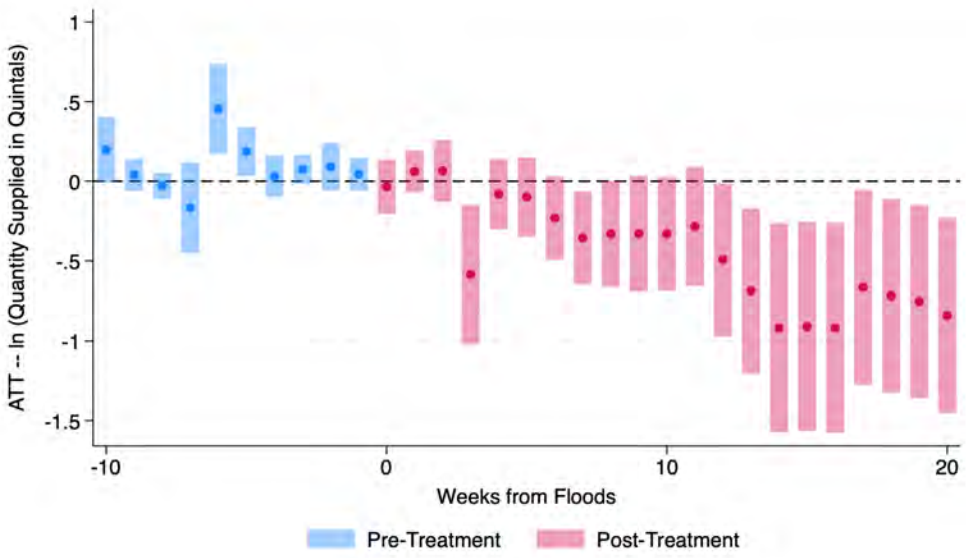

Figure 11: Flood Exposure and Quantity Arrivals in Destination Markets

Notes: Dynamic ATT estimates of flooding on log quantities arriving at destination districts, using a shift-share exposure instrument and event-study specification with crop–destination and week fixed effects, weighted by population. Blue markers and bars denote pre-treatment (weeks –10 to –1), pink markers and bars denote post-treatment (weeks 0 to 20), and shaded areas are 95% confidence intervals.

Figure 12 disaggregates these effects by broad crop category. Fruit arrivals remain flat through week 2 and then begin a steady decline, culminating at –2.5 log points (92%) by week 20—reflecting their high perishability and potentially more localized sourcing. Vegetable volumes similarly hold near baseline for the first few weeks, then fall by about –0.8 log points (55%) from week 3 onward.[13] Grain deliveries, in contrast, exhibit no systematic departure from zero, consistent with their longer shelf life, greater storage capacity, and buffering via state procurement.

[13] While vegetables exhibit a sharper drop in supply at the source level, fruit arrivals fall more steeply at the destination level, reflecting differences in supply chain structure and sourcing patterns. Fruit crops such as apples are often stored at terminal markets by traders using cold storage, allowing for short-term buffering and delayed dispatch to market. However, these crops are frequently sourced from more geographically concentrated regions, making them more vulnerable to sustained upstream disruption once local inventories are exhausted. In contrast, vegetables, despite higher perishability, are typically marketed quickly after harvest and sourced from more spatially dispersed regions, allowing for short-run substitution that may temporarily stabilize arrivals in downstream markets (Asian Development Bank, 2022).

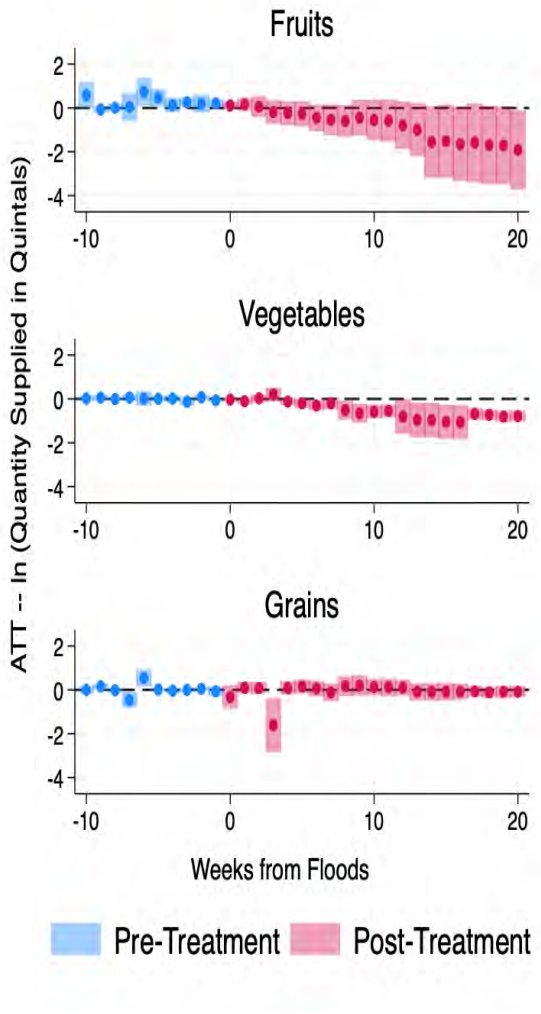

Figure 12: Flood Exposure and Quantity Arrivals, by Crop Type

Notes: ATT estimates of flooding on log quantities arriving at destinations, by crop group: fruits, vegetables, and grains. These are derived using a shift-share exposure instrument and event-study specification with crop–destination and week fixed effects, weighted by population. Blue markers and bars denote pre-treatment (weeks –10 to –1), pink markers and bars denote post-treatment (weeks 0 to 20), and shaded areas are 95% confidence intervals.

These results indicate a lagged supply response: downstream markets only feel the pinch several weeks after upstream floods, with highly perishable fruits and vegetables most severely affected, while grain flows remain resilient. In the next section, we explore how these delayed quantity shortfalls drive adjustments in auction and retail prices.

5.2.2 Auction Prices

Having documented a lagged but substantial collapse in quantities arriving at spot markets, we now turn to how downstream prices adjust in response. Figure 13 presents the event-study ATT estimates of flooding on log auction prices (PKR/kg) for all 16 commodities. Auction prices remain statistically flat through the first six weeks post-flood (weeks 0–6). Beginning around week 7, prices climb steadily, reaching roughly +0.25 log points (28%) by week 10 and peaking near +0.40 log points (49%) by week 20. This lagged response suggests that supply shortfalls in origin markets take several weeks to fully propagate into auction prices downstream.

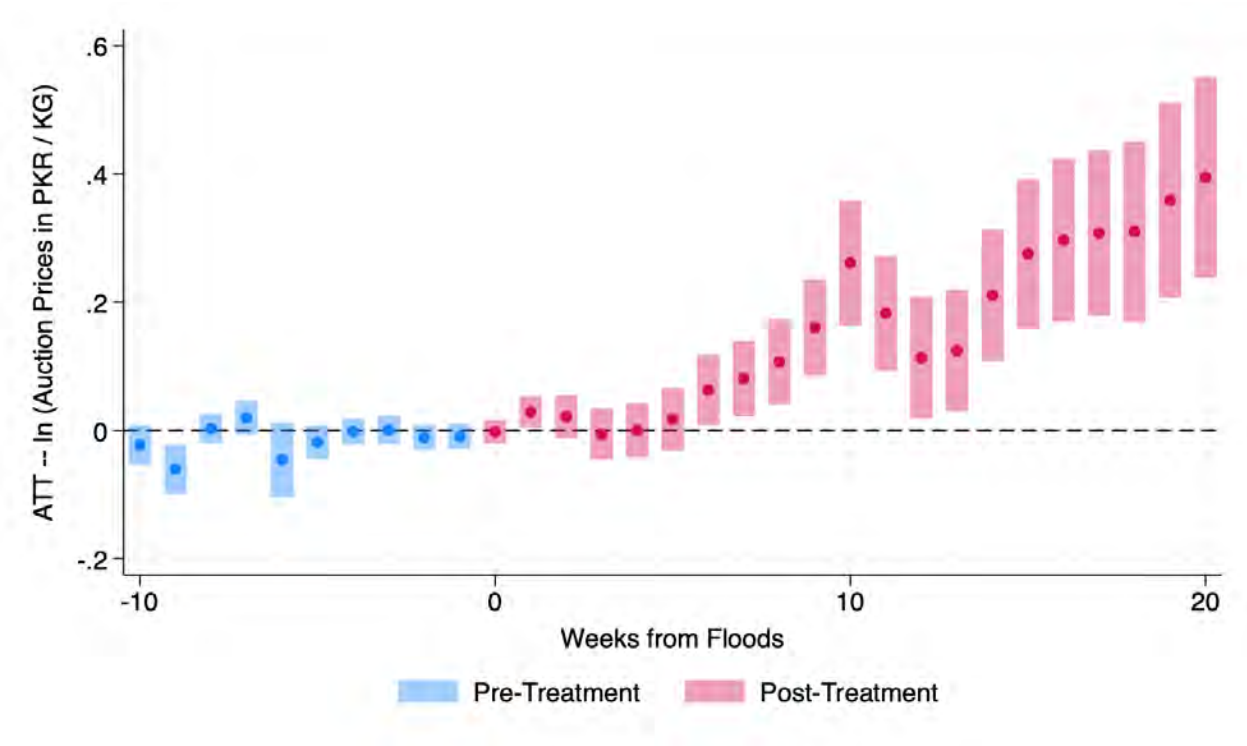

Figure 13: Flood Exposure and Auction Prices

Notes: This figure shows the dynamic impact of source-region flooding on log auction prices across all commodities. Estimates are the ATT from our event-study using a shift–share instrument, with crop–destination and week fixed effects, weighted by destination population. Blue dots and bars (weeks –10 to –1) denote pre-treatment coefficients and 95% CIs; pink dots and bars (weeks 0 to 20) denote post-treatment.

Breaking the auction-price response out by crop type reveals three distinct patterns (Figure 14). Fruits exhibit a muted reaction through week 8 and then a steady climb, rising to about

+0.50 log points—roughly a 65 percent increase—by week 20, mirroring their sustained supply shortfall. Vegetables, by contrast, display an unexpected early-post–flood dip: from roughly weeks 2–7 coefficients turn negative (–0.05 to –0.15 log points, or –5 to –14 percent), before reversing sharply upward after week 8 and peaking near +0.40 log points (49 percent) by week 20. This initial price decline likely reflects a temporary stabilization in downstream markets, possibly due to rerouted shipments out of lower–risk sources or stock releases by intermediaries, before the true magnitude of inundation-induced shortages sets in (Asian Development Bank, 2022). Finally, grain prices remain statistically indistinguishable from zero throughout weeks 0–20, consistent with large-scale state procurement and the lower perishability of staple cereals buffering those markets against short-run flood shocks.

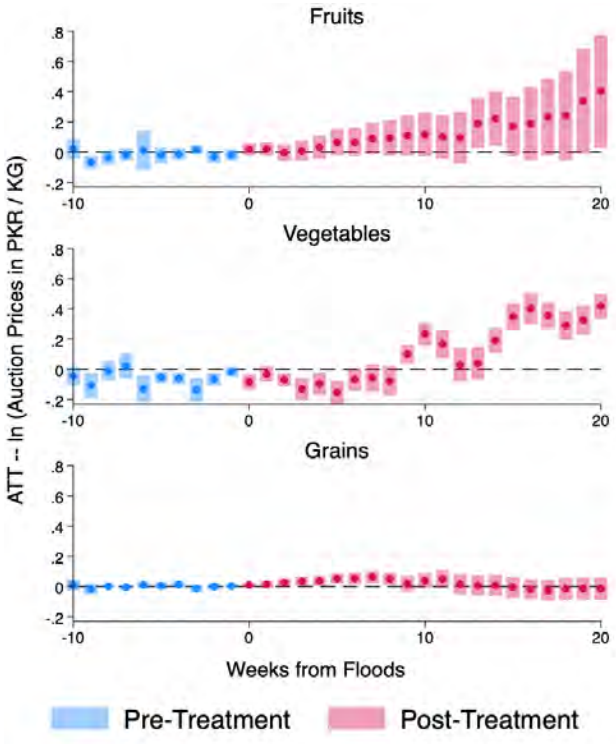

Figure 14: Flood Exposure and Auction Prices, by Crop Type

Notes: This figure shows the average treatment-on-the-treated (ATT) estimates of source-region flooding on log auction prices (PKR/kg), separately for fruits, vegetables, and grains. Estimates come from our shift–share event-study model with crop–destination and week fixed effects, weighted by destination population. Blue markers (weeks –10 to –1) denote pre-treatment estimates; pink markers (weeks 0 to 20) denote post-treatment. Vertical bars are 95% confidence intervals.

5.2.3 Retail Prices and Margins (vegetables)

We next examine how upstream flood-induced price shocks propagate through wholesale markets into retail prices and retailer margins. Figure 15 presents event-study ATT estimates separately for auction and retail prices, focusing on three vegetables (potatoes, onions, tomatoes) across eight major urban districts in Punjab. Auction prices initially remain stable or slightly decline immediately after flooding, reaching their lowest point around week 9. However, from week 10 onward, they increase steadily, ultimately rising to about 0.41 log points (approximately 51%) by week 20. In comparison, retail prices exhibit an earlier and more pronounced response: they start increasing around week 6, rapidly peaking at about 0.50 log points (approximately 65%) by week 10, before moderating somewhat to stabilize around 0.30 log points (approximately 35%) above pre-flood levels.

The differential timing and magnitude of these retail and auction price adjustments imply substantial movements in retailer margins, shown explicitly in Figure 16. Retailer margins, calculated as $\ln((P_{dt}^{\text{Retail}} - P_{dt}^{\text{Auction}})/P_{dt}^{\text{Auction}})$, rise markedly few weeks after flooding onset, peaking at roughly 0.25 log points (approximately 27%) around week 9. This increase suggests that retailers capture a significant portion of scarcity rents arising from upstream shortages. However, margins subsequently begin declining steadily after week 10, gradually approaching zero by week 17, and eventually turning slightly negative by week 20. As shown in the next subsection, this erosion of margins might be due to increased enforcement of price controls.

Taken together, these patterns highlight how initial upstream supply disruptions translate into temporary but substantial gains for retailers, who exploit short-term shortages to expand their margins.

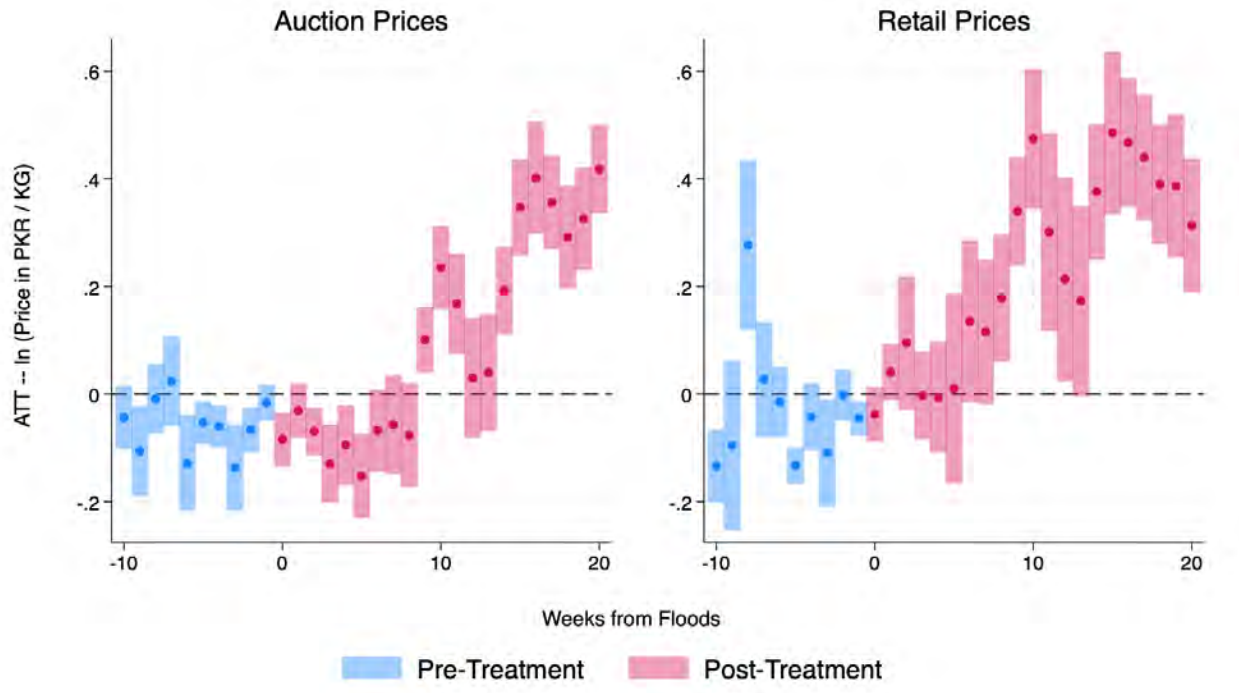

Figure 15: Flood Exposure and Auction vs. Retail Prices

Notes: Event-study ATT estimates for log auction prices (left panel) and log retail prices (right panel) for potatoes, onions, and tomatoes. Estimates are from a shift–share specification with crop–district and week fixed effects, weighted by district population. Blue dots and bars indicate pre-treatment periods (weeks –10 to –1), and pink indicate post-treatment periods (weeks 0 to 20). Shaded areas represent 95% confidence intervals.

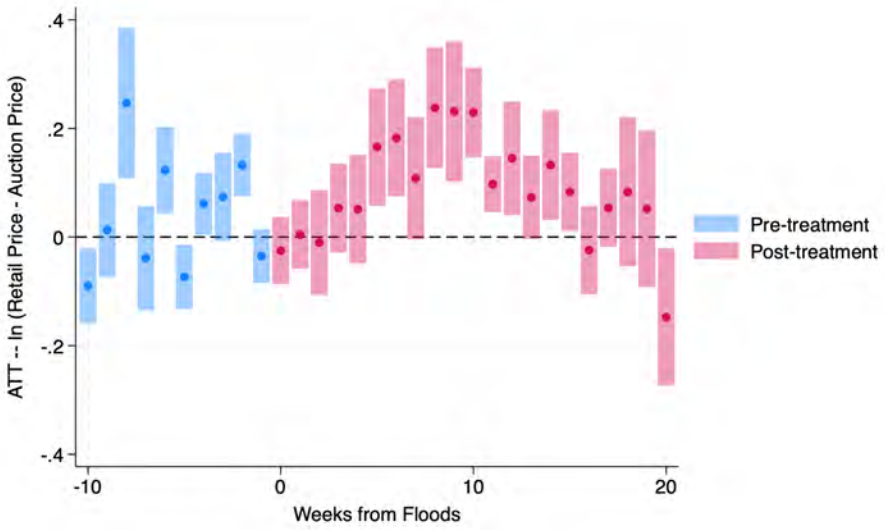

Figure 16: Flood Exposure and Retailer Margins

Notes: Event-study ATT estimates for retailer margins, defined as $\ln((P_{dt}^{\text{Retail}} - P_{dt}^{\text{Auction}})/P_{dt}^{\text{Auction}})$ for potatoes, onions, and tomatoes. Estimates are from a shift–share specification with crop–district and week fixed effects, weighted by district population. Blue dots and bars indicate pre-treatment periods (weeks –10 to –1), and pink indicate post-treatment periods (weeks 0 to 20). Shaded areas represent 95% confidence intervals.

5.2.4 Policy Response (vegetables)

The observed increases in retail prices and retailer margins raise critical questions about the effectiveness and responsiveness of government price control measures. To assess policy responses, we examine how official price ceilings were adjusted following the floods, the extent of retail price overcharging relative to these ceilings, and retailer compliance with official prices.[14]

Figure 17 plots the dynamic adjustments of official price ceilings after the flood shocks. Initially, price ceilings remain largely unchanged, displaying no statistically significant adjustments in the first few weeks following the floods. However, from around week 4 onward, ceilings begin to rise gradually, becoming significantly elevated from week 10 onward. By week 20, ceilings peak at about 0.8 log points (approximately 122%) above pre-flood levels, indicating a substantial policy response to the persistent price pressure in retail markets.

[14] The comparison of raw prices is presented in Figure A3.

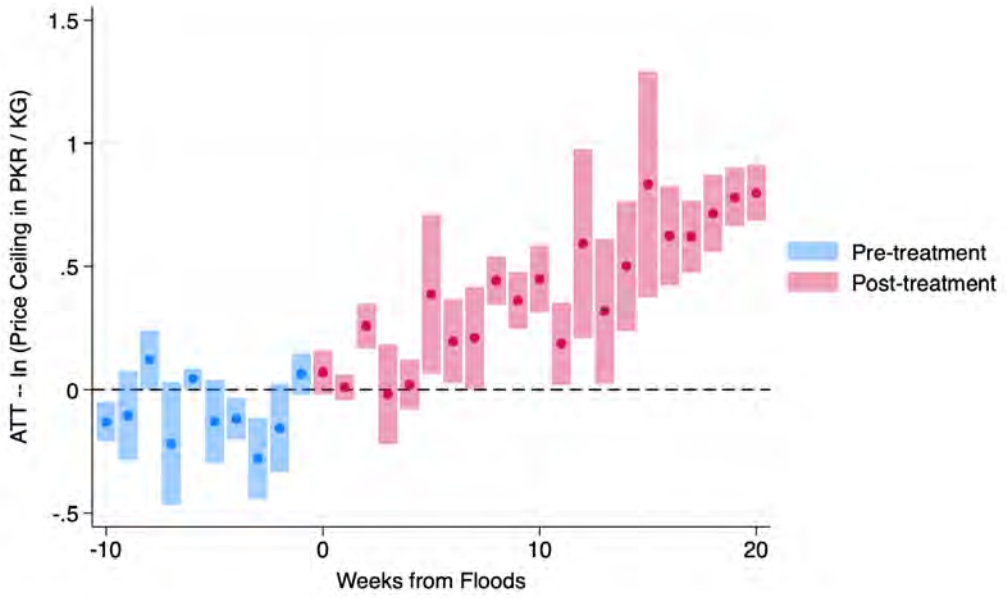

Figure 17: Flood Exposure and Government Price Ceiling Adjustments

Notes: Event-study ATT estimates for the response of official price ceilings (log PKR/kg) to source-region flooding, using a shift–share specification with crop–district and week fixed effects, weighted by district population. Blue dots and bars represent pre-treatment (weeks –10 to –1), and pink dots and bars represent post-treatment (weeks 0 to 20). Vertical bars denote 95% confidence intervals.

Finally, Figure 18 explicitly quantifies retailer compliance, defined as the probability that the official price ceiling exceeds the observed retail price (Compliance $= \mathbf{1}(\text{Ceiling} > \text{Retail})$). Compliance deteriorates significantly post-flood, reaching its lowest point around week 10 (more than 100 percent decrease), indicating widespread noncompliance as ceilings lagged behind market conditions. Subsequently, compliance gradually improves, approaching baseline levels by week 17, likely due to ceiling increases eventually exceeding market prices and enhanced regulatory enforcement.

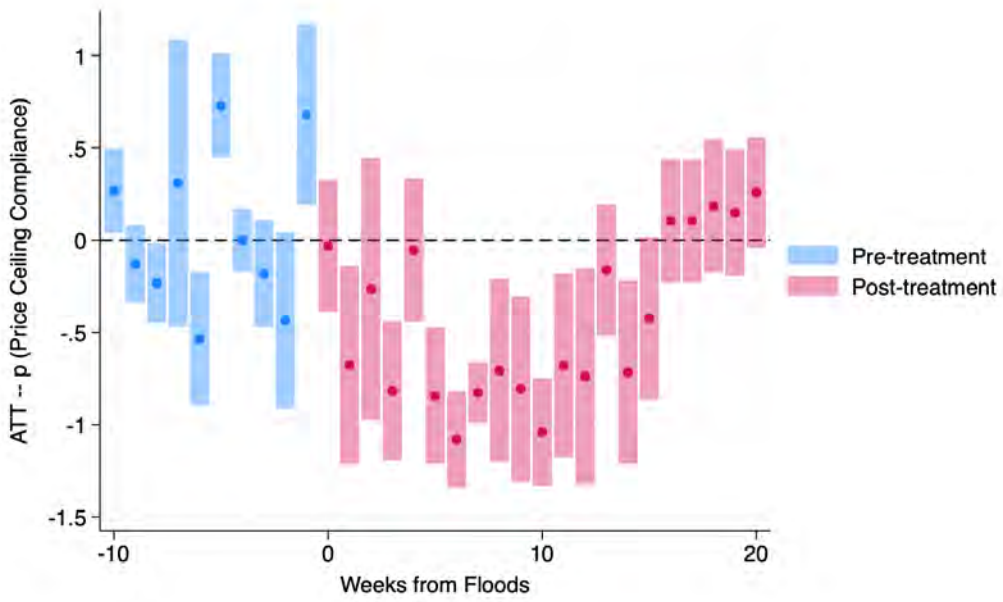

Figure 18: Flood Exposure and Retailer Compliance with Price Ceilings

Notes: Event-study ATT estimates for compliance with official price ceilings, measured as Compliance $= \mathbf{1}(\text{Ceiling} \geq \text{Retail})$. Estimated using a shift–share specification with crop–district and week fixed effects, weighted by district population. Blue dots and bars represent pre-treatment (weeks –10 to –1), and pink dots and bars represent post-treatment (weeks 0 to 20). Vertical bars denote 95% confidence intervals. Negative values indicate reduced compliance.

Collectively, these results underscore that government price interventions following the flood shocks were initially delayed, allowing temporary but substantial noncompliance. Nevertheless, later upward adjustments in price ceilings significantly improved compliance, effectively reducing consumer vulnerability and curbing retail price exploitation in the medium term.

5.3 Road Networks

Before explicitly modeling disruptions to transportation infrastructure, we first document that the impact of source-region flooding varies sharply with the spatial length of supply links. Figure 19 splits source–destination pairs according to whether their average distance lies below or above the median.[15] Two patterns emerge. First, for nearby markets, supply losses are insignificant. Second, for more distant markets, supply collapses sharply at flood onset and remains persistently depressed, with declines exceeding 60 percent and little recovery within the event window. These results indicate that distance itself is a key constraint on supply-chain resilience, motivating a closer examination of whether additional disruptions arise from flooding of transportation networks.

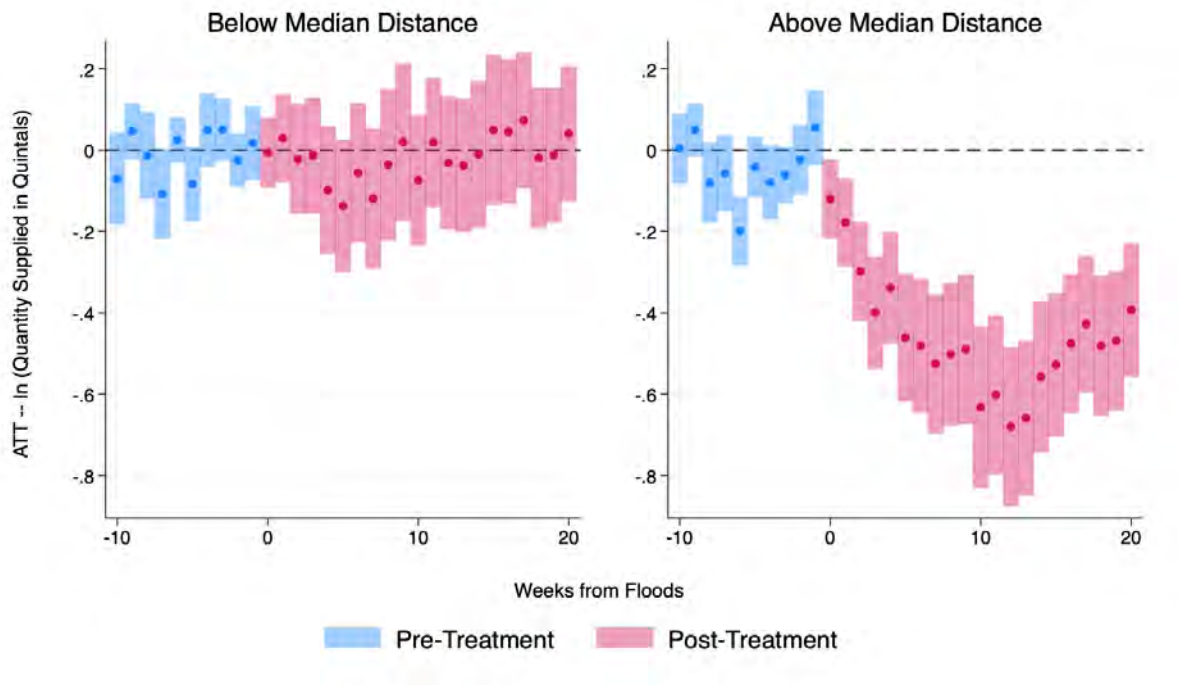

Figure 19: Flood Impacts by Historical Hazard and Source–Destination Distance

Notes: This figure plots event-study estimates of log arrivals from source regions relative to the week of first observed flooding in the source (week 0). Panels split the sample by whether the source–destination distance is below- or above-median. Each specification includes source–destination and week fixed effects; standard errors are clustered at the source region.

[15] Distance is measured as the average pre-flood route distance between each source region and destination market.

6 Conclusion

This study provides novel insights into how extreme climate events, specifically the 2022 floods in Pakistan, disrupt agricultural supply chains in low-income, climate-vulnerable settings. Leveraging high-frequency administrative data linked to satellite flood maps, we trace the temporal and spatial dynamics of this disruption, linking flood exposure to staggered, crop-specific disruptions that persisted well beyond the immediate disaster window.

Using high-frequency administrative data linked to satellite-based flood maps, we deploy a staggered event study framework to estimate how the timing and intensity of inundation affected weekly market arrivals and prices. We find large and persistent post-flood declines in the supply of agricultural commodities, especially for perishable goods such as vegetables and fruits. Staple grains, by contrast, exhibit relative resilience, consistent with greater storability and buffering from government procurement. These supply disruptions were accompanied by sharp increases in both auction and retail prices, as well as widening retailer margins. Notably, these price effects were eventually moderated in districts with stronger enforcement of price ceilings, suggesting that regulatory interventions played a role in shaping market outcomes during the crisis.

These findings contribute to emerging work on how climate shocks propagate through fragmented supply systems. Much of the existing literature emphasizes either micro-level farm outcomes or macroeconomic production shocks; our work highlights the intermediary space of wholesale markets, where both physical flows and pricing dynamics converge. Our findings complement recent work by Balboni et al. (2024), who document disruptions to firm production and trucking networks during the same floods. By focusing on perishable agricultural commodities and public market infrastructure, we extend this literature to food systems and their regulatory governance—key elements for crisis response in low-income settings. The evidence also underscores the role of regulatory institutions: interventions such as procurement or price ceilings can buffer against volatility in the short run, but may also generate unintended spatial asymmetries in availability or affordability.

That said, our analysis has limitations that merit caution. First, while our staggered event study design helps mitigate key confounders by leveraging variation in the timing and intensity of

flood exposure, we cannot fully rule out the influence of concurrent shocks, such as coordinated market behavior, that may affect observed outcomes. Second, the administrative data cover arrivals and auction prices in regulated wholesale markets and do not track final consumer prices and retail margins for fruits and grains. Informal transactions outside the *mandi* system are also not covered by our data. Finally, our commodity categories are broad and do not capture intra-crop variation (e.g., among different varieties or grades), which may mask heterogeneity in perishability or sourcing.

Despite these limitations, the study offers a scalable empirical approach for studying climate disruptions in data-scarce environments. By bridging remote sensing tools with administrative market data, we present a replicable method for tracking climate-driven supply shocks, especially relevant for developing country settings where formal supply chain data remain limited. For policymakers, the results highlight the need to move beyond reactive price controls and toward proactive investments in storage, routing infrastructure, and flood early-warning systems to protect both producers and consumers in increasingly volatile agricultural markets.

References

Ahmad, M., Al Mehedi, M. A., Yazdan, M. M. S., and Kumar, R. (2022). Development of machine learning flood model using artificial neural network (ann) at var river. *Liquids*, 2(3):147–160.

Aker, J. C. (2010). Information from markets near and far: Mobile phones and agricultural markets in niger. *American Economic Journal: Applied Economics*, 2(3):46–59.

Al-Sababhah, N. (2023). Topographic position index to landform classification and spatial planning, using gis, for wadi araba, south west jordan. *Environment and Ecology Research*, 11(1):79–101.

Allen, T. and Arkolakis, C. (2014). Trade and the topography of the spatial economy. *The Quarterly Journal of Economics*, 129(3):1085–1140.

Asian Development Bank (2022). Building horticulture value chains and reducing postharvest losses in pakistan. ADB Brief 235, Asian Development Bank.

Atkin, D. and Donaldson, D. (2015). Who's getting globalized? the size and implications of intranational trade costs. NBER Working Papers 21439, National Bureau of Economic Research, Inc.

Balboni, C., Boehm, J., and Waseem, M. (2024). Firm Adaptation in Production Networks : Evidence from Extreme Weather Events in Pakistan. Technical report. Working paper.

Barrot, J.-N. and Sauvagnat, J. (2016). Input specificity and the propagation of idiosyncratic shocks in production networks *. *The Quarterly Journal of Economics*, 131(3):1543–1592.

Bergquist, L. F. and Dinerstein, M. (2020). Competition and Entry in Agricultural Markets: Experimental Evidence from Kenya†. *American Economic Review*, 110(12):3705–3747.

Boehm, C. E., Flaaen, A., and Pandalai-Nayar, N. (2019). Input linkages and the transmission of shocks: Firm-level evidence from the 2011 tŌhoku earthquake. *The Review of Economics and Statistics*, 101(1):pp. 60–75.

Borusyak, K., Hull, P., and Jaravel, X. (2025). A practical guide to shift-share instruments. *Journal of Economic Perspectives*, 39(1):181–204.

Burgess, R., Deschenes, O., Donaldson, D., and Greenstone, M. (2017). Weather, climate change and death in india. Technical report, University of Chicago.

Casaburi, L. and Reed, T. (2022). Using Individual-Level Randomized Treatment to Learn about Market Structure. *American Economic Journal: Applied Economics*, 14(4):58–90.

Cavallo, A., Cavallo, E., and Rigobon, R. (2014). Prices and supply disruptions during natural disasters. *Review of Income and Wealth*, 60(S2):S449–S471.

Dell, M., Jones, B. F., and Olken, B. A. (2012). Temperature shocks and economic growth: Evidence from the last half-century. *American Economic Journal: Macroeconomics*, 4(3):66–95.

Deschênes, O. and Greenstone, M. (2012). The economic impacts of climate change: Evidence from agricultural output and random fluctuations in weather: Reply. *American Economic Review*, 102(7):3761–3773.

Donaldson, D. (2018). Railroads of the raj: Estimating the impact of transportation infrastructure. *American Economic Review*, 108(4-5):899–934.

Douglas, I. (2009). Climate change, Flooding and Food Security in South Asia. *Food Security*, 1(2):127–136.

Eckstein, D., Kunzel, V., and Schafer, L. (2021). Global Climate Risk Index 2021. Technical report, German Watch.

Farr, T. G., Rosen, P. A., Caro, E., Crippen, R., Duren, R., Hensley, S., Kobrick, M., Paller, M., Rodriguez, E., Roth, L., Seal, D., Shaffer, S., Shimada, J., Umland, J., Werner, M., Oskin, M., Burbank, D., and Alsdorf, D. (2007). The shuttle radar topography mission. *Reviews of Geophysics*, 45(2).

Funk, C., Peterson, P., Landsfeld, M., Pedreros, D., Verdin, J., Shukla, S., Husak, G., Rowland, J., Harrison, L., Hoell, A., and Michaelsen, J. (2015). The climate hazards infrared precipitation with stations—a new environmental record for monitoring extremes. *Scientific Data*, 2(1):150066.

Gorelick, N., Hancher, M., Dixon, M., Ilyushchenko, S., Thau, D., and Moore, R. (2017). Google earth engine: Planetary-scale geospatial analysis for everyone. *Remote Sensing of Environment*, 202:18–27. Big Remotely Sensed Data: tools, applications and experiences.

Hanna, R. and Oliva, P. (2016). Implications of climate change for children in developing countries. *The Future of Children*, pages 115–132.

Headey, D. (2011). Rethinking the global food crisis: The role of trade shocks. *Food Policy*, 36(2):136–146.

Hsiang, S. and Kopp, R. E. (2018). An Economist's Guide to Climate Change Science. *Journal of Economic Perspectives*, 32(4):3–32.

Jensen, R. (2007). The digital provide: Information (technology), market performance, and welfare in the South Indian fisheries sector.

Moazzam, M., Vansarochana, C., and Rahman, A.-U. (2018). Analysis of flood susceptibility and zonation for risk management using frequency ratio model in district charsadda, pakistan. *International Journal of Environment and Geoinformatics*, 5.

Naz, A. and Naz, F. (2024). Issues of Food Pricing Policy in Pakistan and the Way Forward. *The Pakistan Development Review*, 63(3):357–381.

Omid Rahmati, H. Z. and Besharat, M. (2016). Flood hazard zoning in yasooj region, iran, using gis and multi-criteria decision analysis. *Geomatics, Natural Hazards and Risk*, 7(3):1000–1017.

Patel, D. (2025). Floods.

Pekel, J.-F., Cottam, A., Gorelick, N., and Belward, A. S. (2016). High-resolution mapping of global surface water and its long-term changes. *Nature*, 540(7633):418–422.

Pourghasemi, H. R., Amiri, M., Edalat, M., Ahrari, A. H., Panahi, M., Sadhasivam, N., and Lee, S. (2021). Assessment of urban infrastructures exposed to flood using susceptibility map and google earth engine. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 14:1923–1937.

Qamer, F. M., Ahmad, B., Abbas, S., Hussain, A., Salman, A., Muhammad, S., Nawaz, M., Shrestha, S., Iqbal, B., and Thapa, S. (2022). The 2022 pakistan floods: Assessment of crop losses in sindh province using satellite data. Technical report, International Centre for Integrated Mountain Development (ICIMOD); Pakistan Agricultural Research Council's (PARC), Kathmandu, Nepal.

Roy, D., Wulder, M., Loveland, T., C.E., W., Allen, R., Anderson, M., Helder, D., Irons, J., Johnson, D., Kennedy, R., Scambos, T., Schaaf, C., Schott, J., Sheng, Y., Vermote, E., Belward, A., Bindschadler, R., Cohen, W., Gao, F., Hipple, J., Hostert, P., Huntington, J., Justice, C., Kilic, A., Kovalskyy, V., Lee, Z., Lymburner, L., Masek, J., McCorkel, J., Shuai, Y., Trezza, R., Vogelmann, J., Wynne, R., and Zhu, Z. (2014). Landsat-8: Science and product vision for terrestrial global change research. *Remote Sensing of Environment*, 145:154–172.

Rzeszewski, M. (2023). Mapbox. In *Evaluating Participatory Mapping Software*, pages 21–40. Springer.

Sant'Anna, P. H. and Zhao, J. (2020). Doubly robust difference-in-differences estimators. *Journal of Econometrics*, 219(1):101–122.

Stigler, G. J. (1952). *The Theory of Price*. The Macmillan Company, New York, revised edition.

Swain, K. C., Singha, C., and Nayak, L. (2020). Flood susceptibility mapping through the gis-ahp technique using the cloud. *ISPRS International Journal of Geo-Information*, 9(12):720. Submission received: 19 August 2020 / Revised: 20 November 2020 / Accepted: 23 November 2020 / Published: 2 December 2020.

Tariq, A., Yan, J., Ghaffar, B., Qin, S., Mousa, B. G., Sharifi, A., Huq, M. E., and Aslam, M. (2022). Flash flood susceptibility assessment and zonation by integrating analytic hierarchy process and frequency ratio model with diverse spatial data. *Water*, 14(19).

Tehrany, M. S., Pradhan, B., and Jebur, M. N. (2014). Flood susceptibility mapping using a novel ensemble weights-of-evidence and support vector machine models in gis. *Journal of Hydrology*, 512:332–343.

Timmer, C. P. (2010). Reflections on food crises past. *Food Policy*, 35(1):1–11.

United Nations, O. f. t. C. o. H. A. O. (2022a). Pakistan 2022 Floods Response Plan: 01 Sep 2022 - 28 Feb 2023. Technical report, United Nations.

United Nations, U. N. D. P. U. (2022b). Pakistan Floods 2022: Post-Disaster Needs Assessment (PDNA). Technical report, United Nations.

Viktorovic, K. P. (2023). Automation methods in geographic information systems applied to assessing population spatial distribution in cross-border regions.

Appendix A Additional Figures

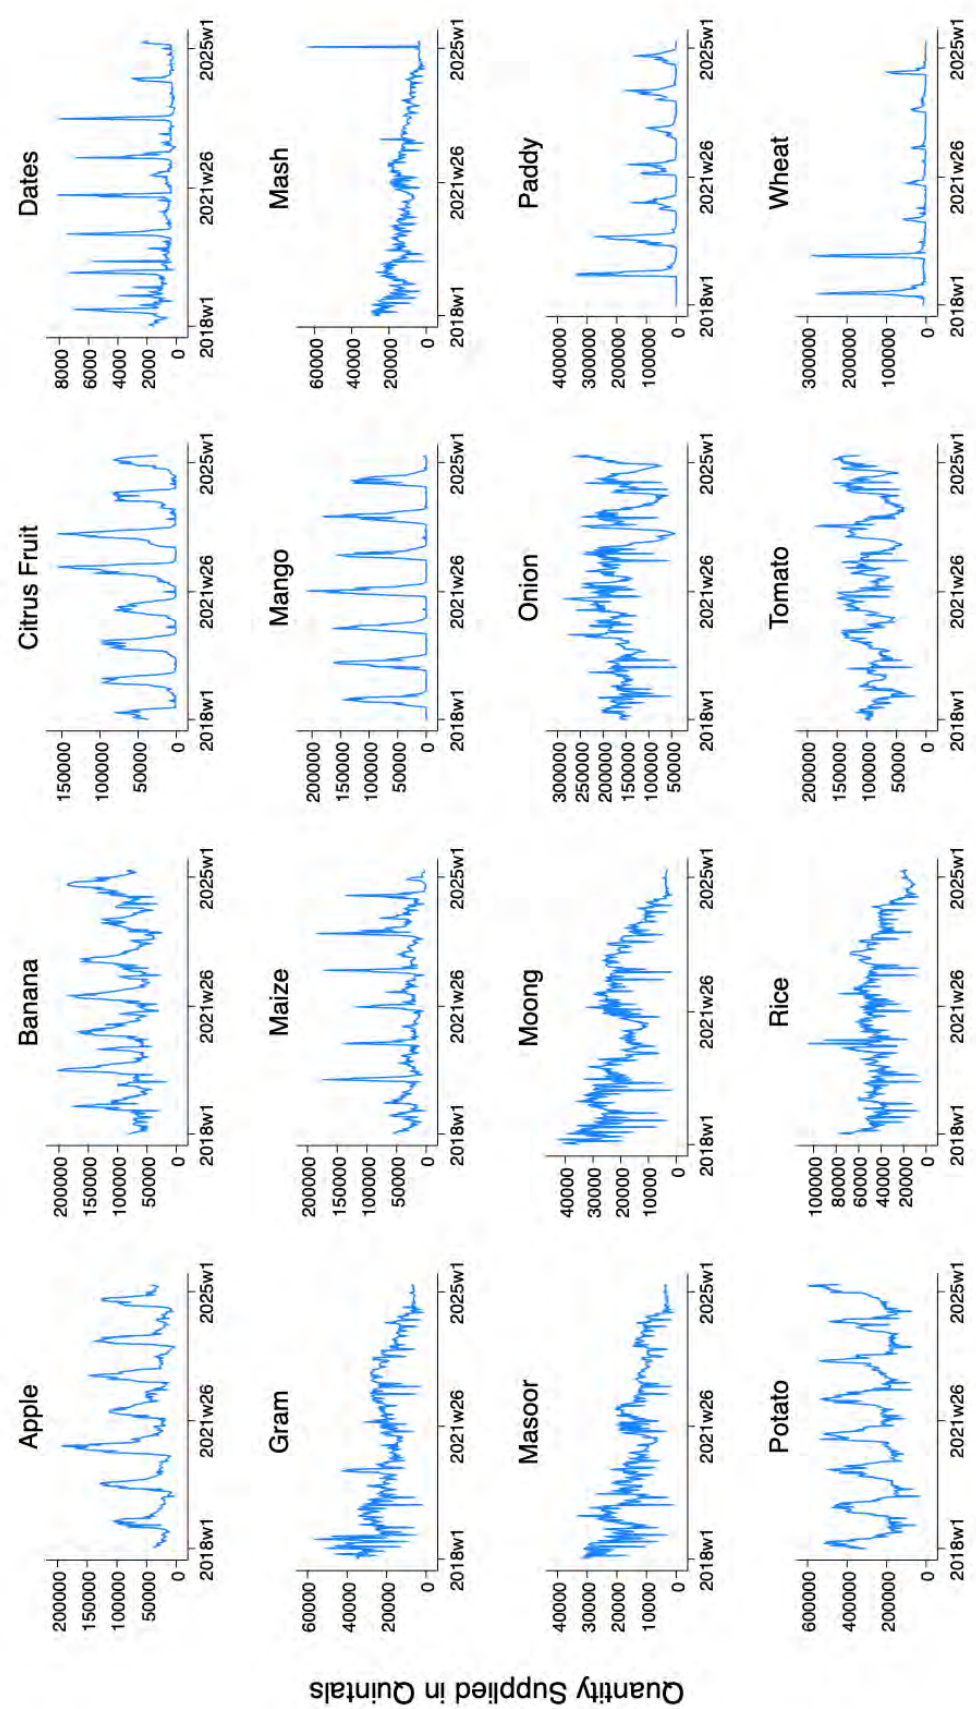

Figure A1: Aggregate Supply, By Crop

Notes: This figure plots weekly aggregate quantities supplied (in quintals) for each crop from 2018w1 to 2025w1. Each panel corresponds to a distinct crop and shows the total quantity arriving across all destination markets from all source regions in a given week.

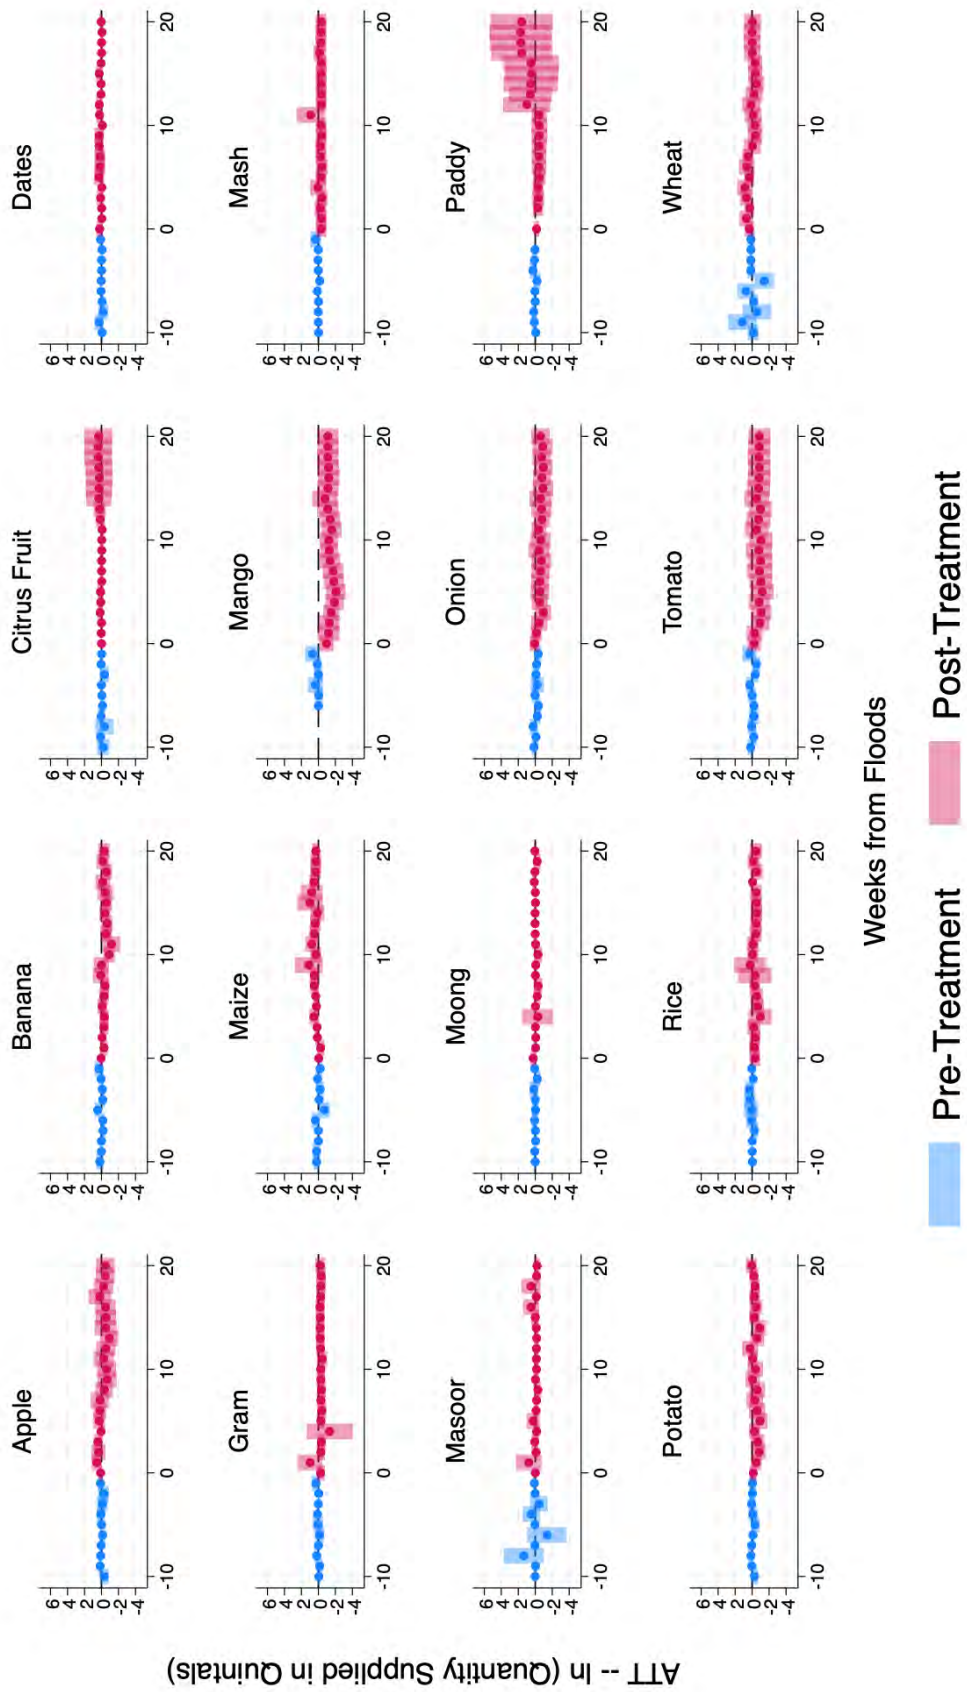

Figure A2: Flood Impact on Quantity Supplied, By Crop

Notes: This figure presents dynamic treatment–on–the–treated (ATT) estimates from our event–study model for each crop separately. The vertical axis shows the ATT in log quintals supplied from source regions, and the horizontal axis measures weeks relative to first observed flooding in a source (week 0). Each panel corresponds to one crop (apples, bananas, citrus fruit, dates, gram, maize, mango, mashed vegetables, masoor, moong, onions, paddy, potatoes, rice, tomatoes, and wheat). Estimates come from a doubly–robust Difference–in–Differences specification with crop–source and week fixed effects, weighted by source area. Blue bars and dots show the pre–treatment period (weeks –10 to –1); pink bars and dots show the post–treatment period (weeks 0 to 20). Shaded bars denote 95 % confidence intervals.

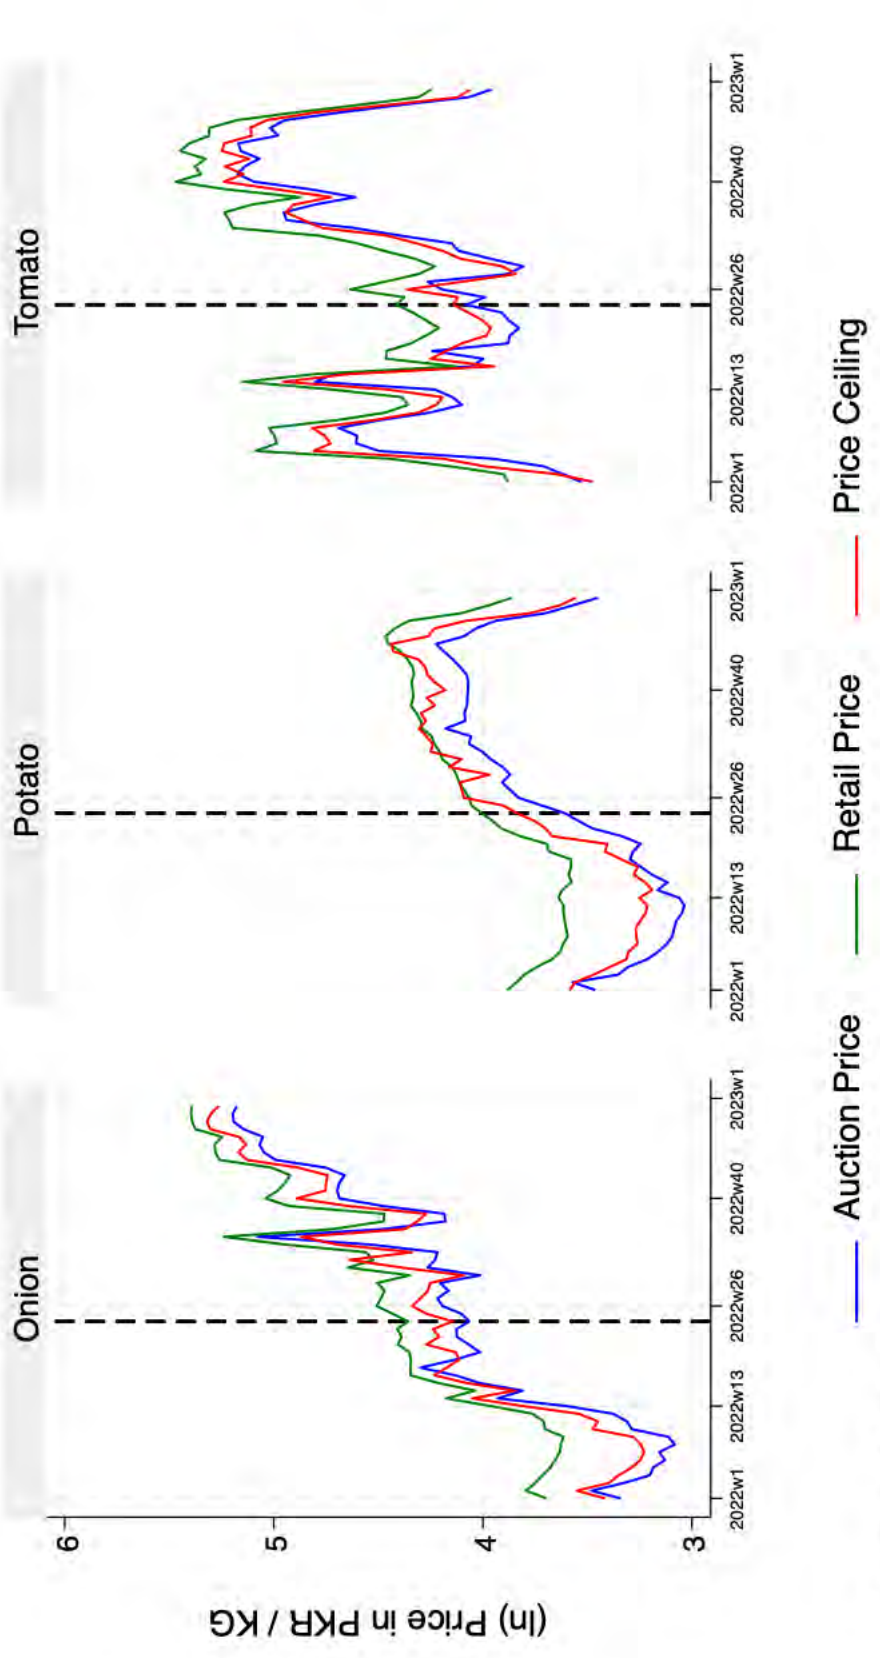

Figure A3: Time series of prices, Vegetables

Notes: This figure plots weekly average auction prices (blue line), retail prices (green line), and official government price ceilings (red line), in log PKR per kilogram, for onions, potatoes, and tomatoes. The vertical dashed line denotes the first widespread flooding event (week 26 of 2022).

Appendix B Flood Detection Algorithm

The calculation of flood water inundation area is done in three steps: 1) Each pixel is assigned a value presented in Table B1; 2) The data are then stitched together and processed using GIS software to remove unwanted pixels, which fall in a range of 0 to 99, and retain the ones that have values exceeding 100). Then flooded area is calculated in m^2 using the formula $\frac{\Sigma_i flood\_frac_i * 140{,}625m^2}{100}$ for all pixels i within the boundary. Note that the area of each pixel is $140,625m^2$ ($375m$ x $375m$). 3) This area is converted into proportion that captures intensity of flooding in the source regions by dividing by the area of the source region.

Assigned Value	**Interpretation**
15	Flood water without water fraction
16/17	Clear Sky/Dry Land
20	Snow
25	Ice
30	Cloud
38	Ice water or mixed water
50	Cloud Water
99	Normal Open Water
100 – 120	Flood Fraction 01 – 20%
121 – 130	Flood Fraction 21 – 30%
131 – 140	Flood Fraction 31 – 40%
141 – 150	Flood Fraction 41 – 20%
151 – 160	Flood Fraction 51 – 60%
161 – 170	Flood Fraction 61 – 70%
171 – 180	Flood Fraction 71 – 80%
181 – 190	Flood Fraction 81 – 90%
111 – 200	Flood Fraction 91 – 100%

Table B1: 8-bit Pixel Assigned Values

Appendix C Flood Risk

There are various methods for estimating flood risk, such as Artificial Neural Networks (ANNs) and other classifiers. For example, ANNs have been used effectively for flood risk prediction due to their ability to model non-linear relationships and handle complex datasets Ahmad et al. (2022); Pourghasemi et al. (2021) However, they require extensive training data and computational resources, and their "black box" nature can make interpretation challenging Tehrany et al. (2014). Boosted regression trees (BRT) and machine learning models have also been applied for urban infrastructure assessment in flood-prone areas, leveraging the computational power of platforms like Google Earth Engine. These models can be prone to overfitting and require careful tuning of parameters Pourghasemi et al. (2021). Similarly, analytical hierarchy process (AHP) and frequency ratio models have been used in conjunction with GIS for flood susceptibility zonation which are relatively easy to implement and allows for the inclusion of expert judgment, these models can be Subjective weighting of criteria and are less effective with large, complex datasets Tariq et al. (2022); Omid Rahmati and Besharat (2016).

Our approach is particularly suited for this study as it avoids the complexities and high resource demands of advanced models by leveraging data for flood risk assessment. Utilizing the built-in functionalities of Google Earth Engine (GEE) for data processing, we ensure a streamlined, scalable, and adaptable workflow. GEE's extensive geospatial datasets and computational power enable comprehensive flood risk assessments by integrating various data sources like satellite imagery, elevation models, and precipitation data Gorelick et al. (2017). The following flowchart illustrates the comprehensive methodology used to assess flood hazard zones:

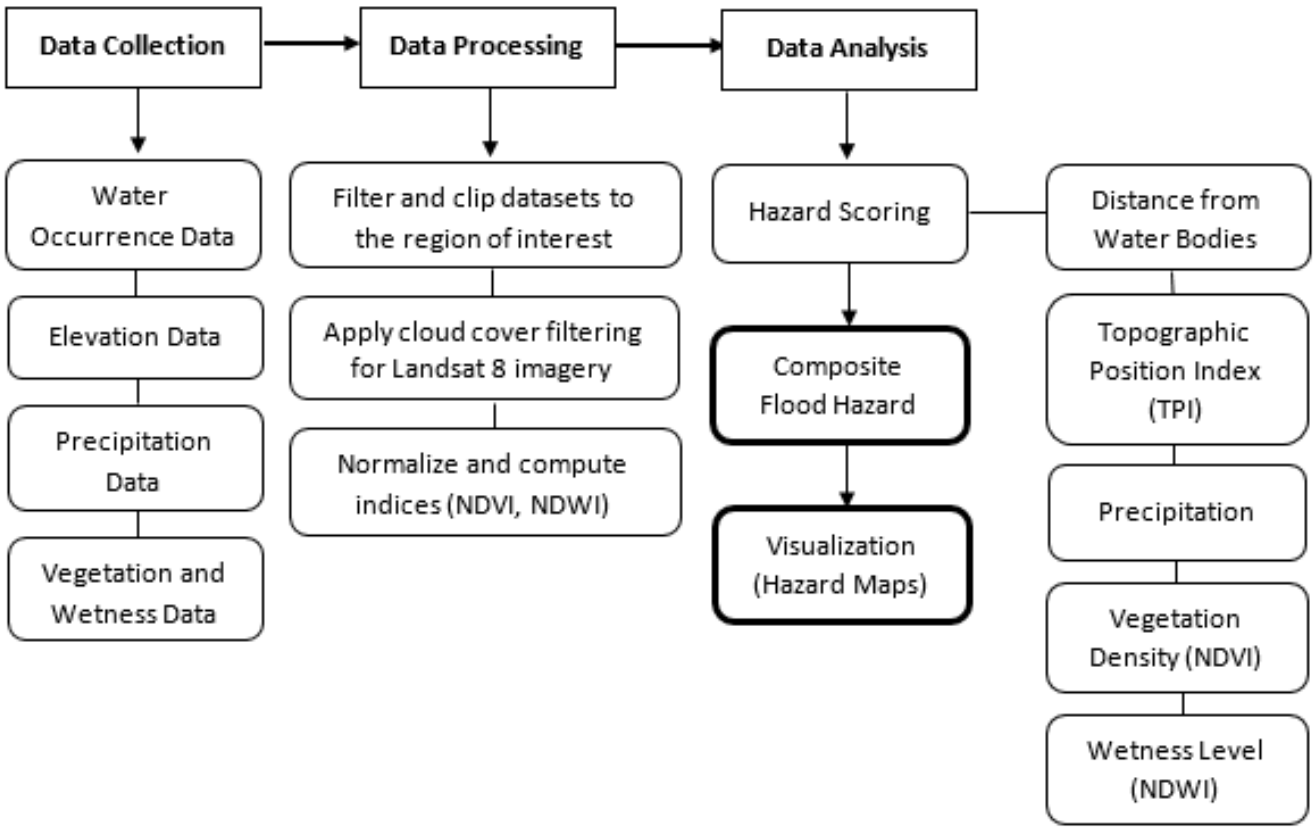

Figure C4: Methodology for Flood Risk Assessment

We use the parameters to generate flood risk or hazard scores, which take integer values from 1 - 5 for each pixel that defines the country spatially. A score of 1 represents areas with the lowest hazard level, while 5 indicates the highest hazard level. We calculate the average risk for each source region by taking the sum of hazard scores for all the pixels spanning that region and dividing by the total number of pixels. When plotted on a map, this exercise generates Figure 5. In this figure, light green regions have low flood risk, while red regions have higher flood risk. We also note that there are isolated red circles scattered on the map. These are areas that surround lakes or local water bodies, and since proximity to permanent water bodies is an important parameter, these regions are assigned higher flood risks. In the middle of the country, passing through D.I. Khan, Layyah, and Muzzaffargarh districts, the path of river Indus is evident as an orange and red strip. Further, the northern and southeastern parts of the country exhibit a high level of risk, whereas the central part of Punjab and the western regions have 'very low' to 'low' risk profiles.

The analysis leverages several key datasets to comprehensively assess flood risk in the study area. The JRC Global Surface Water Mapping Layers (v1.4) is a dataset that contains maps of the location and temporal distribution of surface water from 1984 to 2021. This dataset is crucial for identifying areas prone to water accumulation and flooding by providing statistics on the extent and change of those water surfaces over time Pekel et al. (2016). The dataset is particularly valuable as it includes detailed records of water occurrence, which are used to identify permanent water bodies

and calculate the distance from these bodies to other areas. In our assessment, we utilized this dataset to generate a map of water occurrence, highlighting areas with high water presence, which are likely to be at higher risk of flooding. We then identified permanent water bodies by filtering the dataset to include only areas where water occurrence is greater than 70%. This threshold was chosen to exclude seasonal water bodies and temporary water accumulations, ensuring that only stable, long-term water bodies were considered. The 70% threshold effectively differentiates between permanent water bodies and areas that are occasionally wet due to seasonal variations or temporary floods. Using Google Earth Engine (GEE), we applied the "fastDistanceTransform" function to calculate the distance from these permanent water bodies to all other areas, providing a crucial variable for our flood risk assessment. This distance measure helps determine how likely an area is to experience flooding based on its proximity to permanent water sources. Areas closer to these water bodies are assigned higher hazard scores due to the increased risk of flooding.

The NASA SRTM Digital Elevation 30m (SRTM V3) dataset, produced by the NASA Jet Propulsion Laboratory (JPL), offers digital elevation data at a 30-meter resolution. This elevation data is essential for understanding terrain features and their influence on flood risk, helping to calculate the Topographic Position Index (TPI) and identify flood-prone areas based on elevation and slope Farr et al. (2007), the TPI is derived by subtracting the mean elevation of a specified neighborhood from the elevation at each point, providing insight into whether a point is on a ridge, in a valley, or on flat ground Al-Sababhah (2023), Positive TPI values indicate elevated terrain, while negative values indicate depressions, which are more prone to flooding. In our assessment, the SRTM dataset was clipped to the region of interest and integrated into Google Earth Engine (GEE). Using the elevation data we calculated the TPI by subtracting the focal mean elevation (a moving average of elevation within a defined radius) from the actual elevation. This calculation allowed us to classify the terrain into categories based on flood risk. The TPI values were then used to assign hazard scores: higher scores for areas more susceptible to flooding (depressions) and lower scores for elevated regions (ridges). The TPI was computed using the formula:

$$\text{TPI} = \text{Elevation} - \text{Mean Elevation in a Neighborhood} \tag{7}$$

Climate Hazards Group InfraRed Precipitation with Station data (CHIRPS) is a 30+ year quasi-global rainfall dataset (approximate resolution 0.05°) This high-resolution precipitation data is essential for analyzing rainfall patterns and trends, which are critical for assessing flood risk, as areas with high precipitation are more likely to experience flooding Funk et al. (2015). By incorporating long-term rainfall data, we can better understand the temporal and spatial variability of precipitation, which helps in identifying regions that are more susceptible to flooding due to frequent or intense rainfall events.

The Landsat 8 Imagery dataset, obtained from the United States Geological Survey (USGS), supplies multispectral imagery that includes near-infrared (NIR), red, and green bands. This high-resolution imagery is instrumental in environmental monitoring and analysis due to its ability to capture detailed information about the Earth's surface. Specifically, the NIR, red, and green bands are used to compute indices such as the Normalized Difference Vegetation Index (NDVI) and the Normalized Difference Water Index (NDWI). These indices help assess vegetation density and water content, respectively, which are crucial for identifying areas with high vegetation density and wetness —key factors in flood risk assessment Roy et al. (2014). In our flood risk assessment,

the Landsat 8 imagery was processed to derive the NDVI and NDWI indices. The processing steps included filtering the imagery for cloud cover, selecting the appropriate bands, and clipping the data to the region of interest (ROI) respectively, both indices are calculated using the formula:

$$NDVI : \frac{NIR - RED}{NIR + RED} \tag{8}$$

and

$$NDWI : \frac{GREEN - NIR}{GREEN + NIR} \tag{9}$$

Both NDVI and NDWI were used to assign hazard scores for vegetation density and wetness. Areas with low NDVI values (indicating sparse vegetation) and high NDWI values (indicating high wetness) were assigned higher hazard scores, reflecting their higher flood risk potential integrated at the tehsil level using administrative boundary data, and exported for further analysis and visualization.