When Nature Strikes: Flood Exposure and Import-Based Adaptation Among Indian Firms

Benoît Chaulvet[1]

Université Paris 1 Panthéon Sorbonne
Centre d'Economie de la Sorbonne (CES)

February 2026

Abstract

Extreme floods are expected to intensify under climate change, yet evidence on firm-level adaptation in emerging economies remains limited. This paper estimates the dynamic effects of major flood episodes on Indian manufacturing firms using CMIE ProwessIQ panel data (1990–2006) combined with a measure of flood exposure. I construct treatment on a 5km×5km grid by combining disaster records from EM-DAT and Indian Meteorological Department archives with CHIRPS rainfall anomalies, and link firms to exposed pixels using geocoded locations. Effects are estimated in a stacked event-study design across multiple flood episodes, with treated and control firms matched using several approaches. The results show that floods significantly reduce profit margins among non-importing firms. By contrast, firms that imported before the shock are substantially more resilient. Mechanism tests suggest that importer resilience is not primarily driven by raw-material use, but rather by lower post-shock labor-cost pressure and the preservation of distribution and commercial activity. These findings show that pre-existing access to international input markets can act as a source of firm resilience to climate-related shocks in emerging economies.

Keywords: Natural disasters, Firm resilience, Event study, Panel data
JEL: D22, L6, L7, O12, O53, Q54, Q56, R3

[1] I would like to address my sincere gratitude to Maria Bas, and to Julien Wolfersberger that helped and guided me throughout the process of this paper.
ORCiD: 0009-0003-0134-9642 benoit.chaulvet@etu.univ-paris1.fr

1 Introduction

Climate-related disasters are increasingly central to economic risk, and floods are among the most economically disruptive events in many emerging economies (Intergovernmental Panel on Climate Change, 2021). Climate projections indicate that these risks are likely to intensify. In its Sixth Assessment Report, the IPCC shows that heavy precipitation events that historically occurred once per decade are expected to become more frequent and more intense under higher global warming scenarios. Under a 2°C warming scenario, such events are projected to occur 1.7 times as often and to be 14% more intense; under a 4°C warming scenario, they are projected to occur 2.7 times as often and to be around 30% more intense (Intergovernmental Panel on Climate Change, 2021). These risks are particularly relevant for developing economies, where exposure to climate hazards is often compounded by weaker infrastructure, thinner insurance markets, limited fiscal capacity, and greater dependence on local production networks. Existing evidence shows that the economic and human costs of natural disasters are shaped by income and institutional capacity: poorer countries tend to experience larger losses and slower recovery, while richer and better-governed economies are better able to absorb shocks (Kahn, 2005; Noy, 2009; Loayza et al., 2012).

A growing literature studies how natural disasters affect firm performance in emerging economies (Bas and Paunov, 2024; Nedoncelle and Wolfersberger, 2024). Less is known, however, about which pre-existing firm capabilities help firms withstand these shocks. This paper examines how major floods affect Indian manufacturing firms and whether pre-existing access to international input markets, proxied by importer status, helps firms better withstand these shocks.

India provides a particularly relevant setting to study the economic effects of floods. The country is highly exposed to monsoon-related rainfall shocks: around 80% of annual rainfall is concentrated during the June-September monsoon season, generating large river discharges and recurrent flood risk (Felix, 2025). Floods are also economically costly. Using official flood-damage statistics for 1953-2020, Felix (2025) report average annual flood damages of around \$750 million[2]), with substantial losses to crops, houses, and public infrastructure. Over the same period, floods damaged on average more than 4 million hectares of crops and over 1.2 million houses per year, while causing more than 1,600 deaths annually. These aggregate figures underline the scale of flood exposure in India, but they also mask substantial spatial heterogeneity in where floods occur and which economic units are affected.

This spatial heterogeneity creates an important measurement challenge. Disaster records often identify affected areas at the state or district level, while flood damages are highly localized. Treating all firms in an affected state as equally exposed risks misclassifying both treated and control firms. To address this issue, I construct a high-resolution flood-exposure measure on a 5km×5km grid by combining disaster records from EM-DAT and Indian Meteorological Department archives with CHIRPS rainfall anomalies.

I combine the flood-exposure measure with firm-level data from CMIE ProwessIQ, a panel database maintained by the Centre for Monitoring Indian Economy. ProwessIQ reports detailed financial accounts for Indian companies, including sales, profits, input expenditures, labor costs, and trade-related variables. I link these firm accounts to flood exposure by geocoding firm locations, focusing on manufacturing firms observed between 1990 and 2006. The empirical strategy follows a stacked event-study difference-in-differences design that pools multiple major flood episodes while comparing exposed firms to similar non-exposed firms within each event window (Sun and Abraham, 2021; Cengiz et al., 2019). The preferred specification includes firm, year, and event fixed effects, allowing treatment effects to vary flexibly over event time. To improve comparability between treated and control firms, I construct a matched estimation sample using a hybrid

[2]₹6,429 crore corresponds to ₹64.3 billion rupees. Using an exchange rate of 1 dollar for ₹85 rupees gives us a \$750 million.

Mahalanobis and propensity-score caliper procedure based on pre-shock firm characteristics (Rosenbaum and Rubin, 1983; Diamond and Sekhon, 2013). I then assess robustness using propensity-score matching, inverse probability weighting with trimmed weights, richer controls, state-year fixed effects, alternative clustering, leave-one-event-out tests, and joint pre-trend diagnostics.

The main results show that flood exposure primarily affects firms through profitability rather than through a large and systematic contraction in sales. In the preferred Hybrid specification, floods reduce profit margins among non-importing firms by 3.3 percentage points in the short run and by 2.7 percentage points over the full post-shock window. Firms that imported before the shock display substantially more resilient dynamics: the importer differential is positive and nearly offsets the decline in margins experienced by non-importers. Results on cash-profit levels show a similar short-run pattern, linking the findings to evidence that natural disasters can generate persistent disruptions to enterprise profits and recovery (de Mel et al., 2012). Mechanism tests suggest that importer resilience is not primarily driven by raw-material use, but rather by lower post-shock labor-cost pressure and the preservation of distribution and commercial activity. Sector-level estimates further indicate that the resilience differential is strongest in heavy industry and in rubber, plastics, and other manufacturing sectors. Overall, the results suggest that pre-existing access to international input markets can act as a source of firm resilience by limiting the deterioration in operating profitability after local climate shocks.

This paper contributes to three strands of the literature. First, it adds to research on the economic effects of natural disasters in developing and emerging economies. While existing work has documented aggregate and sectoral effects of disasters (Albala-Bertrand, 1993; Noy, 2009; Loayza et al., 2012), more recent studies examine how firms respond to climate-related shocks (Bas and Paunov, 2024; Nedoncelle and Wolfersberger, 2024). This paper contributes to this literature by providing firm-level evidence on profitability margins and by showing that flood exposure affects profit margins more clearly than sales. Second, it contributes to the measurement of disaster exposure. Because flood damages are highly localized, state-level disaster records may misclassify firm exposure. I address this by constructing a high-resolution flood treatment that combines disaster records with rainfall anomalies on a 5km×5km grid, allowing exposure to vary within affected states. Third, the paper connects the literature on climate-shock resilience with the literature on firms' international linkages. By showing that pre-shock importer status predicts stronger post-flood resilience, the analysis highlights international input access as a potential adaptation margin through which firms in emerging economies can better withstand localized climate shocks.

The remainder of the paper proceeds as follows. Section 2 outlines the mechanisms linking floods to firm performance and discusses the role of importing from a theoretical standpoint. Section 3 describes the data and the construction of the pixel-level flood treatment. Section 4 presents the stacked event-study design and the estimation strategy. Section 5 reports the main results and diagnostics. Sections 7 and 8 discuss the findings and conclude. Section 6 aims to test the robustness of the results. Appendix 9 provides additional robustness checks, balance diagnostics, and event-time coefficient tables.

2 Conceptual Framework

This section outlines the economic mechanisms through which floods may affect firm performance and explains why importer status may shape firms' resilience to these shocks. The framework is not intended as a formal model, but rather as a guide for the empirical analysis. It clarifies why flood exposure may affect profitability even in the absence of large sales losses, why pre-shock access to international input markets may matter, and which adjustment margins can be examined in the data.

2.1 Floods and firm profitability

Floods can affect firm performance through several channels. The most direct channel is physical disruption: inundation can damage plants, machinery, inventories, and local infrastructure, reducing productive capacity and delaying production. Floods can also disrupt transport networks, prevent workers from reaching production sites, and affect local demand by damaging households, businesses, and public infrastructure. These disruptions imply that flood exposure may reduce output, sales, and profits, consistent with evidence that natural disasters can generate persistent disruptions to enterprise recovery and firm performance (de Mel et al., 2012; Bas and Paunov, 2024; Nedoncelle and Wolfersberger, 2024).

However, the effects of floods need not appear only through sales. A firm may continue selling after a flood while facing higher operating costs, lower productivity, or more difficult input and labor-market conditions. Recent evidence on climate shocks and firms shows that heavy rainfall can reduce markups and revenue productivity through production efficiency losses and higher marginal costs (Bas and Paunov, 2024). Similarly, disaster-induced disruptions may slow enterprise recovery and affect profit dynamics when firms lack the resources needed to rebuild or adjust (de Mel et al., 2012). Floods can also disrupt suppliers, transport networks, and local inventories, raising input and logistics costs or delaying production. These mechanisms are consistent with evidence that shocks propagate through production networks when inputs are specific or difficult to substitute (Barrot and Sauvagnat, 2016; Carvalho et al., 2021; Boehm et al., 2019). In this case, sales may remain partly preserved, but profit margins deteriorate because production and distribution become more costly.

This distinction is important for interpreting firm resilience. If floods mainly reduce sales, resilience reflects the ability to maintain output and demand. If floods primarily reduce profit margins, resilience instead reflects the ability to absorb local disruptions without a sharp increase in costs relative to revenues. The empirical analysis therefore considers both sales and profitability outcomes, including profit margins and cash-profit levels.

2.2 Importing as a resilience margin

Importer status may capture several pre-existing capabilities that help firms withstand local flood shocks. The most direct channel is input-sourcing flexibility. If floods disrupt domestic suppliers, transport infrastructure, or local inventories, firms with access to foreign input markets may be better able to substitute away from locally constrained supply chains and maintain production. This mechanism is closely related to evidence that shocks propagate through production networks when inputs are specific and difficult to replace (Barrot and Sauvagnat, 2016; Carvalho et al., 2021; Boehm et al., 2019).

Importing may also proxy for a broader set of firm capabilities. Firms engaged in international trade tend to differ systematically from purely domestic firms in size, productivity, capital intensity, and organizational capacity (Bernard et al., 2007). Access to imported inputs can expand the set of available varieties, raise productivity, and increase firms' flexibility in adjusting production (Amiti and Konings, 2007; Goldberg et al., 2010). In the Indian context, Goldberg et al. (2010) show that access to new imported intermediate inputs contributed to domestic product growth, suggesting that imported inputs can shape firms' production possibilities. Importer status may therefore capture not only direct access to foreign suppliers, but also stronger logistics networks, more diversified supplier relationships, better access to working capital, and greater managerial capacity. These broader capabilities are consistent with evidence linking financial constraints to participation in international trade (Manova, 2013) and with evidence from Indian manufacturing that management practices affect productivity, quality, efficiency, and inventory management (Bloom et al., 2013).

The empirical analysis examines several margins through which importer resilience may operate. I first test whether flood exposure affects sales, cash-profit levels, and profit margins. I then study whether importer resilience is associated with changes in raw-material use, labor-cost pressure, and distribution or commercial activity. Some potentially relevant channels, such as insurance coverage, supplier contracts, inventory buffers, credit-line access, and the use of trade credit, are not directly observed in the data. I therefore interpret these channels as complementary mechanisms and use observed accounting margins to assess which forms of adjustment are most consistent with the data (de Mel et al., 2012; Lai et al., 2022).

2.3 Empirical predictions

The previous discussion leads to three empirical predictions. First, flood exposure should reduce firm performance. If floods primarily disrupt production or demand, the effect should appear through lower sales and cash profits. If firms are able to maintain sales while facing higher production, labor, or distribution costs, the effect should appear more strongly through lower profit margins. The empirical analysis therefore examines both sales and profitability outcomes.

Second, firms with pre-shock access to international input markets should be more resilient to local flood shocks. Importer status may allow firms to draw on broader supplier networks, reorganize sourcing, or rely on organizational and logistical capabilities that are less dependent on local conditions. I therefore test whether firms that imported before the shock experience smaller post-flood declines than otherwise similar non-importing firms.

Third, importer resilience should be visible in firm adjustment margins. If the main channel is input-sourcing flexibility, importers should be better able to maintain raw-material use or cash-profit levels after floods. If importer status instead captures broader organizational and commercial resilience, differences should appear through lower labor-cost pressure or the preservation of distribution and commercial activity. The mechanism analysis therefore examines raw-material use, labor-cost intensity, distribution expenses, and related accounting margins to identify which channels are most consistent with the observed resilience pattern.

3 Data and Flood Exposure

This section describes the firm-level data, the construction of the flood-exposure measure, and the base firm-year sample used in the analysis. The empirical unit in the underlying panel is a firm–year. Flood exposure varies by firm location and event year, and the stacked event-study sample used for estimation is constructed from this base panel.

3.1 Firm-level data and outcomes

Firm-level outcomes and covariates are drawn from CMIE ProwessIQ, a panel database maintained by the Centre for Monitoring Indian Economy. ProwessIQ compiles harmonized firm financial accounts from audited annual reports for listed firms and a broad set of unlisted firms. The analysis focuses on geocoded manufacturing firms observed between 1990 and 2006, which provides several years of pre- and post-event information around the major flood episodes used in the empirical analysis.

I restrict the sample to firms with usable geographic information and non-missing values for the main accounting variables used in the analysis. Firms are linked to flood exposure using their geocoded locations, which are assigned to the corresponding 5 km×5 km rainfall pixel. Firms without usable location information are excluded from the analysis sample.

The main outcome is the firm's profit margin, defined as cash profits divided by sales. This outcome captures the extent to which flood exposure affects operating profitability relative to revenues. I also examine three additional outcomes. First, I use IHS-transformed cash profits to study whether the results also appear in profit levels. Second, I use log sales to test whether flood exposure leads to a contraction in firm revenues. Third, I use log raw-material expenditures as an input-side outcome, motivated by the possibility that floods disrupt production and input sourcing. Together, these outcomes make it possible to distinguish between revenue losses, profit-level effects, and changes in margins.

The main heterogeneity variable is pre-shock importer status. For each flood episode, I define a firm as an importer if it reports importing before the shock. This pre-shock definition avoids mechanically capturing post-disaster adjustments in importing behavior and allows importer status to be interpreted as a pre-existing firm capability. I also use firm-level accounting variables such as assets, leverage, age, productivity proxies, labor costs, raw-material expenditures, and distribution expenses to construct matched samples, examine balance, and study mechanisms.

3.2 Flood events and exposure construction

To identify major flood episodes, I combine disaster records from EM-DAT with event information from Indian Meteorological Department archives. The empirical analysis focuses on the major flood years retained in the baseline event stack: 1993, 1994, 1995, and 2000. For each event year, I identify the states documented as affected by major flooding and combine this information with high-resolution precipitation data.

The precipitation data come from CHIRPS, which provides gridded rainfall information at approximately 5 km resolution. I aggregate rainfall to the monthly level and construct standardized rainfall anomalies for each pixel and calendar month. For pixel p, month m, and event year e, the anomaly is defined as:

$$A_{p,m,e} = \frac{R_{p,m,e} - \overline{R}_{p,m}}{\sigma\left(R_{p,m}\right)}. \tag{1}$$

Here, $R_{p,m,e}$ is monthly rainfall in the event year, while $\overline{R}_{p,m}$ and $\sigma\left(R_{p,m}\right)$ are the long-run mean and standard deviation for the same pixel and calendar month.

Flood exposure is then defined on a 5 km × 5 km grid. For each event year, I identify the peak flood month using IMD narratives and classify a pixel as exposed if it lies in an affected state and its rainfall anomaly during the event month exceeds the baseline threshold. Formally, pixel-level exposure is:

$$G_{p,e} = \mathbf{1}\left[\text{state}(p) \in \mathcal{S}_e \text{ and } A_{p,m(e),e} > \tau\right]. \tag{2}$$

Here, $\mathcal{S}_e$ is the set of affected states in event year e, $m(e)$ is the peak flood month, and τ is the anomaly threshold. In the baseline specification, $\tau = 1$, corresponding to rainfall at least one standard deviation above the long-run pixel-month mean.

Figure 1 illustrates the spatial information underlying the treatment definition for the July 1993 flood episode. The figure overlays documented flooded states with high-resolution precipitation anomalies, showing why a purely state-level treatment would be too coarse: rainfall intensity varies substantially within affected states. The map is intended as a visual illustration of the rainfall variation used to construct the exposure measure; the formal treatment variable is based on the standardized anomaly defined above. Equivalent maps for the other flood episodes are reported in the appendix.

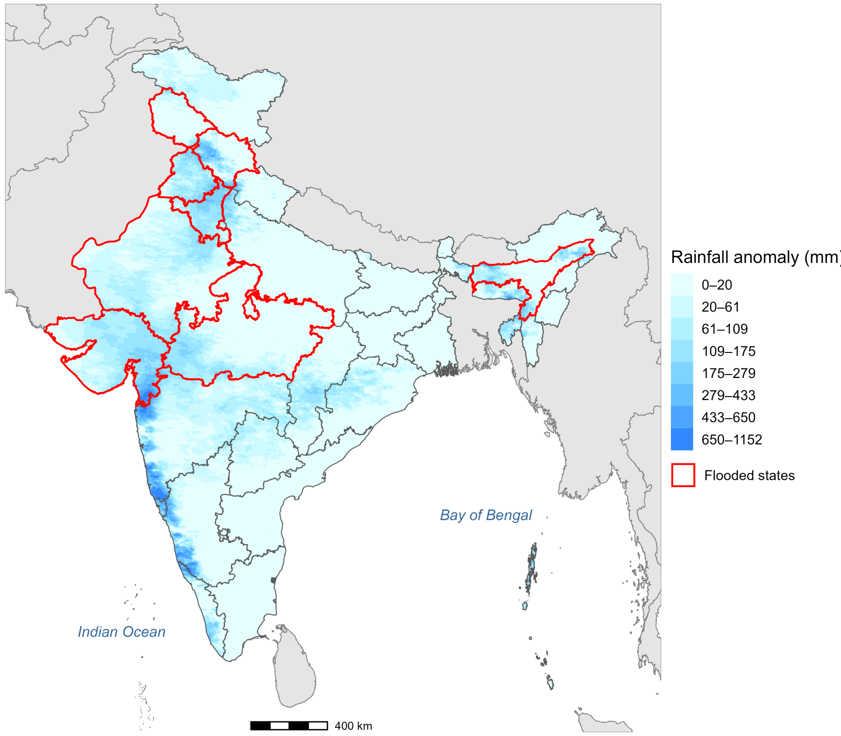

Figure 1: **Precipitation anomalies and documented flooded states in July 1993**

*Notes:*The figure illustrates the spatial inputs used to construct the flood-exposure measure for the July 1993 flood episode. Red outlines indicate states identified as flooded in disaster records, while the raster layer shows positive CHIRPS rainfall anomalies relative to the pixel-specific July mean, measured in millimeters. The formal treatment variable combines documented flood exposure at the state level with the standardized pixel-level rainfall anomaly defined in the text. Equivalent maps for the remaining flood episodes are shown in the appendix.

This procedure addresses a central measurement challenge. Disaster records often identify affected areas at the state or district level, while flood exposure is spatially localized. Treating all firms in an affected state as exposed would risk misclassifying firms located far from the most intense rainfall anomalies. By combining documented flood states with pixel-level rainfall anomalies, the treatment definition restricts attention to major flood episodes while allowing exposure to vary within affected states.

I assign each firm to a pixel using geocoded firm locations and define firm-level exposure in event e as $G_{i,e} = G_{p(i),e}$, where $p(i)$ denotes the pixel containing firm i. This event-specific treatment variable is then used to construct the stacked event-study dataset described in Section 4.

3.3 Base analysis sample and descriptive statistics

Table 1 reports summary statistics for the base firm-year analysis sample before estimator-specific matching or weighting restrictions. The sample contains 41,356 firm-year observations for 6,622 geocoded manufacturing firms across 27 states. Firms in the sample are relatively large and internationally connected, consistent with the coverage of the ProwessIQ database: 67.9% of firm-year observations correspond to importers and 57.0% to exporters.

The average cash-profit margin is close to zero, but the median is 4.8%, indicating substantial dispersion in profitability across firm-years. The median firm is 17 years old, while the higher mean age reflects the presence of older firms in the sample. Estimator-specific support restrictions for the Hybrid, PS matching, and IPW specifications are discussed in Section 4 and reported in the appendix.

Table 1: Summary statistics: base firm-year sample

Variable	Observations	Mean	Std. dev.	Median
Cash-profit margin	41,356	0.001	0.272	0.048
IHS cash profits	41,356	1.112	1.938	1.029
Log sales	41,356	3.732	1.705	3.743
Log raw materials	41,356	3.029	1.848	3.137
Importer	41,356	0.679	0.467	1.000
Exporter	41,356	0.570	0.495	1.000
Firm age	37,980	25.700	60.800	17.000
Log assets	41,356	3.520	1.520	3.370
Leverage	41,356	0.386	0.302	0.360
TFP	41,356	1.140	0.638	1.095

Notes: The table reports summary statistics for the base firm-year analysis sample before estimator-specific matching or weighting restrictions. The sample includes geocoded manufacturing firms observed between 1990 and 2006. Cash-profit margin is defined as cash profits divided by sales and winsorized at the 1st and 99th percentiles. IHS denotes the inverse hyperbolic sine transformation. Importer and exporter are binary variables, so their means correspond to sample shares.

4 Empirical Strategy

This section describes the stacked event-study design used to estimate the dynamic effects of flood exposure on firm outcomes. I first explain how event-specific panels are constructed and stacked. I then present the estimating equation and clarify the interpretation of the domestic/non-importer effect, the importer differential, and the implied importer effect. Finally, I describe the matched and weighted samples, document covariate balance and pre-trend diagnostics, and discuss inference.

4.1 Stacked event-study design

The empirical strategy exploits variation in firm exposure to major flood episodes across space and time. For each flood event $e \in 1993, 1994, 1995, 2000$, I construct an event-specific panel of firms observed before and after the flood year. Event time is defined as $k = t - e$, where t is the calendar year and e is the flood-event year. The treatment indicator $G_{i,e}$ equals one if firm i is located in a flood-exposed pixel for event e, and zero otherwise. Firms located in non-exposed pixels within the same event stack form the event-specific control group.

I then pool the event-specific panels into a stacked event-study dataset. This approach follows recent applications of stacked difference-in-differences designs and is well suited to settings with multiple treatment episodes and event-specific comparison groups (Cengiz et al., 2019; Wing et al., 2024). It allows the analysis to combine several major flood episodes while keeping treatment timing aligned relative to each event. It also allows each flood episode to have its own comparison group and event fixed effect. Because the same firm-year can fall within the event window of more than one flood episode, observations in the stacked dataset are indexed by firm, event, and calendar year. The stacked design therefore compares exposed and non-exposed firms within event-specific windows rather than relying on a single aggregate treatment timing.

The main analysis focuses on dynamic treatment effects around the flood year and reports both event-time coefficients and average effects over post-treatment windows. I summarize post-flood effects over three windows: the short-run window $[0, 2]$, the later post-shock window $[3, 4]$, and the full post-shock window $[0, 4]$. The year immediately before the flood, $k = -1$, is omitted and used as the reference period.

4.2 Estimating equation and treatment-effect interpretation

The baseline specification estimates the effect of flood exposure separately for pre-shock non-importers and for pre-shock importers. This event-study specification is motivated by the recent difference-in-differences literature emphasizing the interpretation of dynamic treatment effects under treatment-effect heterogeneity (Sun and Abraham, 2021; Callaway and Sant'Anna, 2021; Borusyak et al., 2024). Let $Y_{i,e,t}$ denote an outcome for firm i in event stack e and calendar year t. Let $D^k_{i,e,t}$ be an indicator equal to one when firm i is in event time k in event stack e and is exposed to the flood:

$$D^k_{i,e,t} = \mathbf{1}\{t - e = k\} \times G_{i,e}. \tag{3}$$

The baseline event-study equation is:

$$Y_{i,e,t} = \sum_{k \neq -1} \beta_k D^k_{i,e,t} + \sum_{k \neq -1} \delta_k \left(D^k_{i,e,t} \times Importer^{pre}_{i,e} \right) + \alpha_i + \lambda_t + \mu_e + \varepsilon_{i,e,t}. \tag{4}$$

The specification includes firm fixed effects α_i, calendar-year fixed effects λ_t, and event fixed effects μ_e. Firm fixed effects absorb time-invariant differences across firms, calendar-year fixed effects absorb aggregate shocks common to firms in a given year, and event fixed effects account for level differences across flood-event stacks.

The coefficients β_k measure the dynamic effect of flood exposure for treated firms that were not importers before the shock. The coefficients δ_k measure the differential effect for firms that were importers before the shock. The implied effect for treated importers at event time k is therefore $\beta_k + \delta_k$. A positive δ_k indicates that pre-shock importers are more resilient than otherwise comparable non-importers after flood exposure.

For readability, the main tables report average treatment effects over post-treatment windows rather than every event-time coefficient. For a post-treatment window $\mathcal{K}$, the domestic/non-importer effect is the average of the relevant β_k coefficients:

$$\overline{\beta}_{\mathcal{K}} = \frac{1}{|\mathcal{K}|} \sum_{k \in \mathcal{K}} \beta_k. \tag{5}$$

while the importer differential is:

$$\overline{\delta}_{\mathcal{K}} = \frac{1}{|\mathcal{K}|} \sum_{k \in \mathcal{K}} \delta_k. \tag{6}$$

The implied importer effect is $\overline{\beta}_{\mathcal{K}} + \overline{\delta}_{\mathcal{K}}$. Standard errors for these window-level effects are computed using the variance-covariance matrix of the estimated event-study coefficients.

4.3 Construction of matched and weighted samples

A central concern is that firms located in flood-exposed pixels may differ systematically from non-exposed firms. To improve comparability between exposed and non-exposed firms, I use matching and weighting methods based on pre-shock firm characteristics. This follows the logic of propensity-score adjustment and matching estimators, which aim to compare treated and control observations with similar observable characteristics (Rosenbaum and Rubin, 1983; Abadie and Imbens, 2006; Stuart, 2010). Since the preferred specification estimates regressions after matching, the analysis is also related to recent work on post-matching inference (Abadie and Spiess, 2022).

The preferred estimator uses a Hybrid matching procedure that combines Mahalanobis matching with a propensity-score caliper. Matching is performed within event-specific samples using pre-shock firm characteristics. The matching variables include baseline firm size, profitability, leverage, age, productivity proxies, and pre-shock international exposure. This procedure restricts the comparison to treated and control firms with similar pre-shock characteristics and common support.

The Hybrid matched sample is the preferred specification because it balances flexibility and comparability. Mahalanobis matching preserves closeness in the multidimensional covariate space, while the propensity-score caliper prevents matches between firms with very different predicted probabilities of exposure. The main regressions are then estimated on the matched stacked sample using the corresponding matching weights.

I also estimate two alternative specifications. The first uses propensity-score matching alone. The second uses inverse probability weighting (IPW), with weights trimmed to reduce the influence of extreme observations (Hirano et al., 2003). These alternative estimators provide robustness checks against the possibility that the results are specific to the preferred Hybrid matching procedure. Since each estimator imposes different support restrictions, the sample sizes differ across specifications. Appendix Table 9 reports estimator-specific support statistics, and Appendix Table 10 reports the corresponding model sample sizes for the main cash-profit-margin specification.

Before outcome-specific estimation restrictions, the preferred Hybrid stacked sample contains 11,683 observations, 980 firms, and 1,539 firm-event observations. In the main cash-profit-margin estimating equation, the Hybrid specification uses 11,479 observations. The corresponding main-model samples contain 12,603 observations under PS matching and 63,116 observations under IPW.

4.4 Identification, balance, pre-trends, and inference

The identifying assumption is that, conditional on the matched or weighted comparison design and fixed effects, exposed and non-exposed firms would have followed parallel outcome trends in the absence of flood exposure. In this setting, the assumption requires that pre-shock differences between firms located in flood-exposed pixels and comparable firms located in non-exposed pixels do not predict differential post-shock changes in outcomes, after controlling for firm, year, and event fixed effects.

I provide several pieces of supporting evidence for this assumption. First, I examine covariate balance before and after matching or weighting. The balance diagnostics compare exposed and non-exposed firms on pre-shock firm characteristics used in the matching and weighting procedures, including firm size, assets, leverage, age, productivity proxies, pre-shock international exposure, and geographic coordinates. This follows the logic of matching and propensity-score methods, which use pre-treatment observables to improve comparability between treated and control units (Rosenbaum and Rubin, 1983; Abadie and Imbens, 2006; Stuart, 2010; Hirano et al., 2003). I also report extended balance checks on pre-treatment outcomes and mechanism variables, including profitability, input-cost ratios, and distribution-related expenses, in Appendix Table 8. These diagnostics assess whether the Hybrid matching procedure improves comparability between treated and control firms on observables. Table 2 shows that the preferred Hybrid estimator provides the strongest balance improvement: the mean absolute standardized difference falls from 0.227 to 0.058, and 90% of covariates have adjusted standardized differences below 0.10.

Table 2: Covariate balance across estimators

Estimator	Covariates	Mean SMD before	Mean SMD after	Max SMD after	Share below 0.10 after
Hybrid	10	0.227	0.058	0.225	0.900
PS matching	10	0.227	0.128	0.264	0.400
IPW	10	0.227	0.111	0.220	0.500

Notes: The table reports covariate balance diagnostics for the frozen estimation samples. SMD denotes the absolute standardized mean difference between exposed and non-exposed firms. The *before* column reports balance before matching or weighting, while the *after* columns report balance after applying the corresponding estimator-specific weights. The balance covariates are pre-treatment log sales, assets, TFP, leverage, electricity use, firm age, importer status, exporter status, latitude, and longitude. A common rule of thumb is that standardized mean differences below 0.10 indicate good balance. Hybrid refers to Mahalanobis matching with a propensity-score caliper; PS refers to propensity-score matching; IPW refers to inverse probability weighting with weights capped at the 97.5th percentile.

Figure 2 reports covariate balance for the preferred Hybrid estimator. The figure shows that the Hybrid procedure substantially reduces imbalance for most pre-treatment covariates. The main remaining imbalance concerns longitude. This is not surprising in this setting, since flood exposure is spatial by construction and treated firms need not be located in the same east–west part of India as comparison firms. I therefore interpret this residual imbalance transparently and assess robustness to specifications with richer controls and state-year fixed effects.

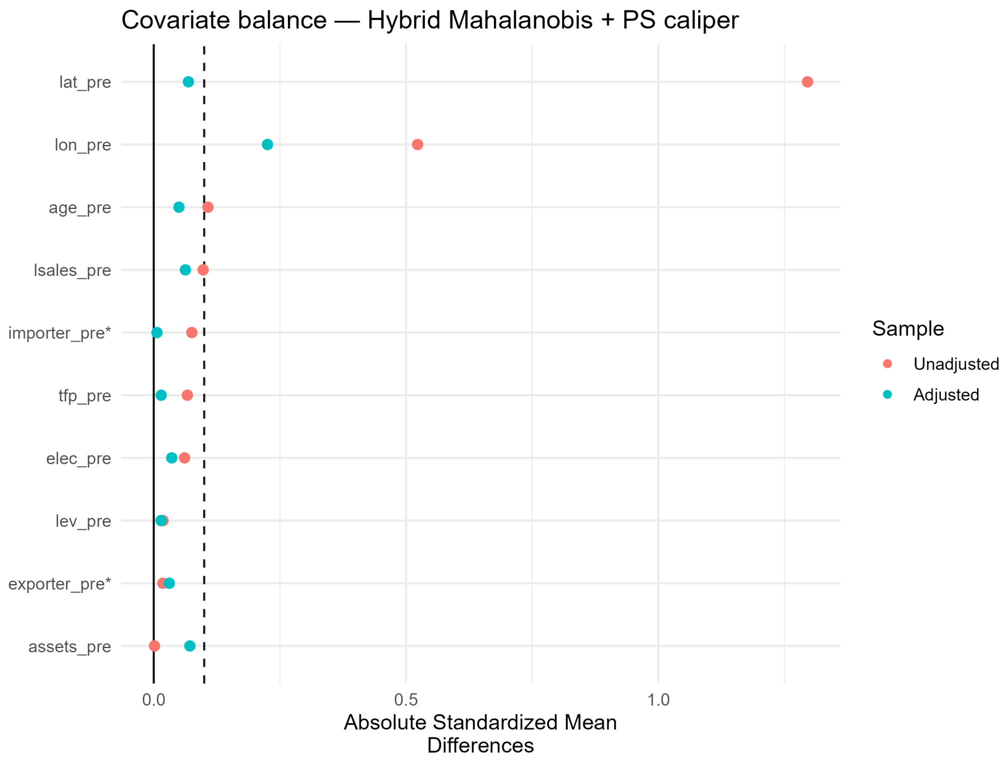

Figure 2: **Covariate balance for the preferred Hybrid estimator**

Notes: The figure reports absolute standardized mean differences between flood-exposed and non-exposed firms before and after applying the Hybrid matching weights. The Hybrid estimator combines Mahalanobis matching with a propensity-score caliper. Balance is assessed using pre-treatment covariates measured over the pre-shock window: log sales, assets, TFP, leverage, electricity use, firm age, importer status, exporter status, latitude, and longitude. The dashed vertical line marks the conventional 0.10 threshold for standardized mean differences. Asterisks indicate binary covariates as reported by the balance-plot routine.

I also report extended balance checks on pre-treatment outcomes and mechanism variables, including profitability, input-cost ratios, and distribution-related expenses, in Appendix Table 8. These extended diagnostics show that Hybrid again performs best among the three estimators, reducing the mean absolute standardized difference from 0.113 to 0.084. Some residual imbalance remains for profit-level and distribution-related variables, so these checks are interpreted as transparency diagnostics rather than as additional identifying assumptions.

Second, the event-study design allows the parallel-trends assumption to be assessed directly through pre-treatment coefficients. The coefficients for event times before the flood test whether exposed and non-exposed firms were already on differential trajectories before the shock. Appendix Table 13 reports joint Wald tests for the baseline non-importer treatment leads and the importer differential leads over the window $[-5, -2]$. For the main outcome, cash-profit margin, the tests do not reject the null of no differential pre-trends across all estimators. The same is true for log sales and log raw materials. The evidence is weaker for IHS-transformed cash profits, where some specifications reject or are close to rejecting the null for importer differential leads. I therefore interpret the profit-level results more cautiously and place greater weight on the cash-profit-margin results, which show stronger pre-trend support.

Third, I assess the robustness of the estimates to alternative specifications. The appendix reports results using alternative estimators, richer controls, state-year fixed effects, alternative clustering choices, leave-one-event-out specifications, and alternative sample restrictions. These checks examine whether the main results are driven by a particular estimator, event, or inference choice.

Inference allows for correlation in shocks within firms over time and across firms within states. Standard errors are clustered by firm and state in the main specifications, following the logic of multiway clustered inference (Cameron et al., 2011). This two-way clustering accounts for serial correlation at the firm level and spatially correlated shocks at the state level. Robustness checks using alternative clustering choices are reported in the appendix.

5 Results

This section presents the main empirical results. I first estimate the effect of flood exposure on firm profitability and test whether pre-shock importers are more resilient than otherwise comparable non-importers. I then examine whether importer status captures a distinct resilience margin relative to other pre-shock firm characteristics. Finally, I study potential mechanisms by asking whether importers preserve margins through differences in labor-cost intensity and distribution-related activity.

The results show a consistent pattern. Flood exposure reduces profitability among firms that were not importers before the shock. By contrast, pre-shock importers experience a significantly smaller decline in profit margins. In the preferred Hybrid specification, the short-run effect of flood exposure on the cash-profit margin of non-importers is negative, while the importer differential is positive and of similar magnitude. This implies that importers are largely insulated from the profitability losses experienced by otherwise comparable non-importing firms. The results are strongest for cash-profit margins, remain directionally similar for cash-profit levels, and are weaker for sales and raw-material use.

Figure 3 plots the event-study coefficients for the cash-profit margin in the preferred Hybrid specification. The figure shows the dynamic version of the main result: after flood exposure, non-importers experience a decline in profitability, while the importer differential offsets most of this loss. Table 3 summarizes these dynamics over the post-treatment windows $[0, 2]$ and $[0, 4]$.

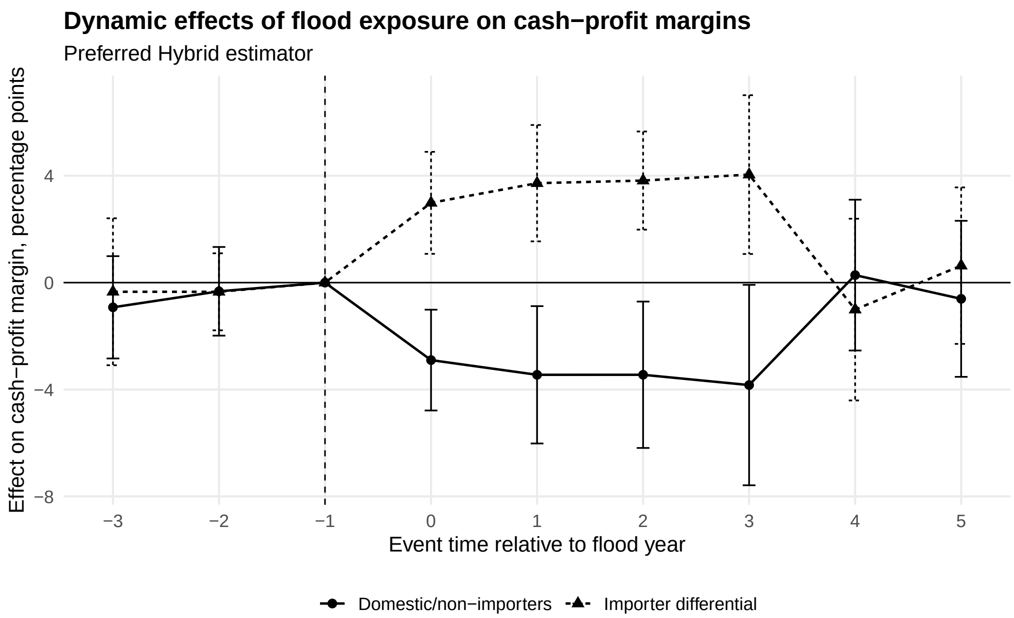

Figure 3: **Dynamic effects of flood exposure on cash-profit margins**

Notes: The figure plots event-study coefficients from the preferred Hybrid specification for the cash-profit-margin outcome. The omitted reference period is $k = -1$. The domestic/non-importer series reports the effect of flood exposure for firms without pre-shock import links. The importer differential series reports the additional effect for firms with pre-shock import links. Effects are reported in percentage points. Confidence intervals are based on standard errors clustered by firm and state.

5.1 Main effects on profitability

Table 3 reports the main results for the preferred Hybrid estimator. The table presents average post-flood effects for four firm outcomes: cash-profit margins, IHS-transformed cash profits, log sales, and log raw-material use. For each outcome, I report the effect for domestic non-importers, the additional differential effect for pre-shock importers, and the implied net effect for importers.

Table 3: Main effects on firm outcomes

	Cash-profit margin	Asinh cash profits	Log sales	Log raw materials
Post-treatment window [0,2]				
Domestic/non-importers	-0.033***	-0.196***	-0.031	-0.023
	(0.010)	(0.044)	(0.029)	(0.040)
Importer differential	0.035***	0.213***	0.049**	0.052
	(0.008)	(0.035)	(0.021)	(0.032)
Implied importer effect	0.002	0.018	0.018	0.029
Post-treatment window [0,4]				
Domestic/non-importers	-0.027***	-0.122**	-0.008	0.000
	(0.010)	(0.059)	(0.039)	(0.059)
Importer differential	0.027***	0.103*	0.024	0.032
	(0.008)	(0.056)	(0.028)	(0.045)
Implied importer effect	0.000	-0.019	0.016	0.032
Observations	11,479	11,479	11,479	11,479
R^2	0.414	0.573	0.913	0.885
Within R^2	0.004	0.003	0.004	0.004
Fixed effects	Firm, year, event	Firm, year, event	Firm, year, event	Firm, year, event
Estimator	Hybrid	Hybrid	Hybrid	Hybrid

Notes: The table reports average event-study effects for the preferred Hybrid estimator. The domestic/non-importer effect is the average post-flood coefficient for firms without pre-shock import links. The importer differential is the additional effect for pre-shock importers. The implied importer effect is the sum of both coefficients. Standard errors clustered by firm and state are reported in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

The magnitude of the cash-profit-margin effect is economically meaningful. In the short-run window $[0, 2]$, flood exposure reduces the cash-profit margin of non-importing firms by 3.3 percentage points. The importer differential is 3.5 percentage points, which almost exactly offsets this loss. Over the full post-shock window $[0, 4]$, the non-importer loss remains sizable at 2.7 percentage points, while the importer differential again offsets nearly the entire decline. The implied effect for pre-shock importers is therefore close to zero in both post-treatment windows.

The remaining outcomes clarify the nature of this effect. The results for IHS-transformed cash profits are directionally similar: non-importers experience a decline in profit levels, while the importer differential is positive, especially in the short run. By contrast, the effects on log sales and log raw-material use are smaller and less precisely estimated. This suggests that importer resilience is not primarily driven by higher post-flood output or input expansion, but by the preservation of margins.

This result is related to the literature on shock propagation through production and supply-chain linkages, but it emphasizes a different margin. Barrot and Sauvagnat show that supplier-specific input relationships can transmit idiosyncratic shocks to downstream firms (Barrot and Sauvagnat, 2016). Boehm, Flaaen, and Pandalai-Nayar use the 2011 Tohoku earthquake to show that input linkages can transmit a natural-disaster shock across firms and countries (Boehm et al., 2019). Carvalho et al. similarly show that the Tohoku shock propagated through firm-level supply chains (Carvalho et al., 2021). In contrast, the estimates here suggest that pre-existing import links can also operate as a buffer against localized domestic shocks. The importer differential does not reflect higher post-flood sales; instead, it reflects a smaller deterioration in profit margins. This interpretation is consistent with work showing that imported intermediates can improve firm performance by expanding input choice, quality, or sourcing flexibility (Amiti and Konings, 2007; László Halpern, 2015).

5.2 Alternative resilience characteristics

A remaining concern is that pre-shock importer status may proxy for broader firm strength rather than for access to international input networks. This concern is important because the international-trade literature shows that globally engaged firms tend to differ systematically from purely domestic firms in size, productivity, and market orientation (Bernard et al., 2007). To assess this possibility, Table 4 compares importer status with other pre-shock resilience characteristics, including exporter status, firm size, assets, TFP, leverage, and age. Importer status remains the most robust predictor of post-flood resilience in the short run. This suggests that the main result is not simply driven by larger, more productive, or internationally oriented firms in general. Age also predicts stronger resilience over the full post-shock window, which is consistent with the idea that older firms may have accumulated organizational capacity or more stable business relationships. However, the short-run resilience pattern is most clearly associated with pre-shock importer status. These estimates should be interpreted as diagnostic comparisons rather than as separate causal effects of each firm characteristic.

Table 4: Alternative pre-shock resilience characteristics

	Importer	Exporter	Sales	Assets	TFP	Leverage	Age
Short-run impact [0,2]							
	0.025***	0.008	0.002	0.000	0.001	0.000	0.004
	(0.007)	(0.009)	(0.006)	(0.001)	(0.005)	(0.005)	(0.003)
Full post-shock average [0,4]							
	0.016***	0.011	0.001	-0.001	0.003	0.000	0.014***
	(0.005)	(0.010)	(0.005)	(0.001)	(0.004)	(0.005)	(0.003)
Observations	11,479	11,479	11,479	11,479	11,479	11,479	11,479
R^2	0.420	0.420	0.420	0.420	0.420	0.420	0.420
Within R^2	0.013	0.013	0.013	0.013	0.013	0.013	0.013
Firms Fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year Fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Event Fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: The dependent variable is the cash-profit margin. Each column reports the average post-flood interaction between treatment exposure and a pre-shock resilience characteristic. Continuous pre-shock variables are standardized. Standard errors clustered by firm and state are reported in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

5.3 Mechanisms

The main results indicate that importers are more resilient in terms of profitability, but they do not by themselves identify the channel. Table 5 examines whether the importer differential is associated with differences in labor-cost intensity, distribution-related activity, sales, and wages after flood exposure.

The mechanism estimates help explain why the profitability effect is concentrated in margins rather than in sales. In the short run, flood exposure increases the wage-to-sales ratio of non-importers by 1.3 percentage points, while the importer differential is -1.7 percentage points. This implies that importers do not experience the same increase in labor-cost intensity after floods. Similarly, distribution expenses fall for non-importers, while the importer differential is positive and offsets most of this decline. These patterns suggest that non-importers experience both margin compression and a contraction in commercial or distribution activity, whereas importers are better able to preserve operating capacity.

Figure 4 summarizes the two main mechanism patterns. For non-importers, flood exposure increases labor-cost intensity and reduces distribution-related activity. For importers, the differential effects move in the opposite direction, offsetting these changes. This supports the interpretation that importers' profitability

Table 5: Mechanisms: labor-cost intensity and distribution capacity

	Cash-profit margin	Wage-to-sales ratio	Log distribution expenses	Log sales	Log wages
Post-treatment window [0,2]					
Domestic/non-importers	-0.033***	0.013***	-0.043**	-0.031	-0.015
	(0.010)	(0.005)	(0.018)	(0.029)	(0.016)
Importer differential	0.035***	-0.017***	0.039**	0.049**	0.026
	(0.008)	(0.006)	(0.016)	(0.021)	(0.022)
Implied importer effect	0.002	-0.004	-0.004	0.018	0.011
Post-treatment window [0,4]					
Domestic/non-importers	-0.027***	0.010**	-0.061***	-0.008	-0.015
	(0.010)	(0.005)	(0.020)	(0.039)	(0.020)
Importer differential	0.027***	-0.016***	0.062***	0.024	0.030
	(0.008)	(0.006)	(0.014)	(0.028)	(0.024)
Implied importer effect	0.000	-0.006	0.001	0.016	0.015
Observations	11,479	11,479	11,479	11,479	11,479
R^2	0.414	0.799	0.885	0.913	0.962
Within R^2	0.004	0.006	0.005	0.004	0.006
Fixed effects	Firm, year, event	Firm, year, event	Firm, year, event	Firm, year, event	Firm, year, event
Estimator	Hybrid	Hybrid	Hybrid	Hybrid	Hybrid

Notes: The table reports mechanism outcomes using the preferred Hybrid estimator. The wage-to-sales ratio captures labor-cost intensity. Distribu capture commercial/distribution activity. The distribution channel should be interpreted together with the rich-control pre-trend diagnostics disc Standard errors clustered by firm and state are reported in parentheses. $^*p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

advantage reflects both lower margin compression and greater preservation of commercial operating capacity after floods.

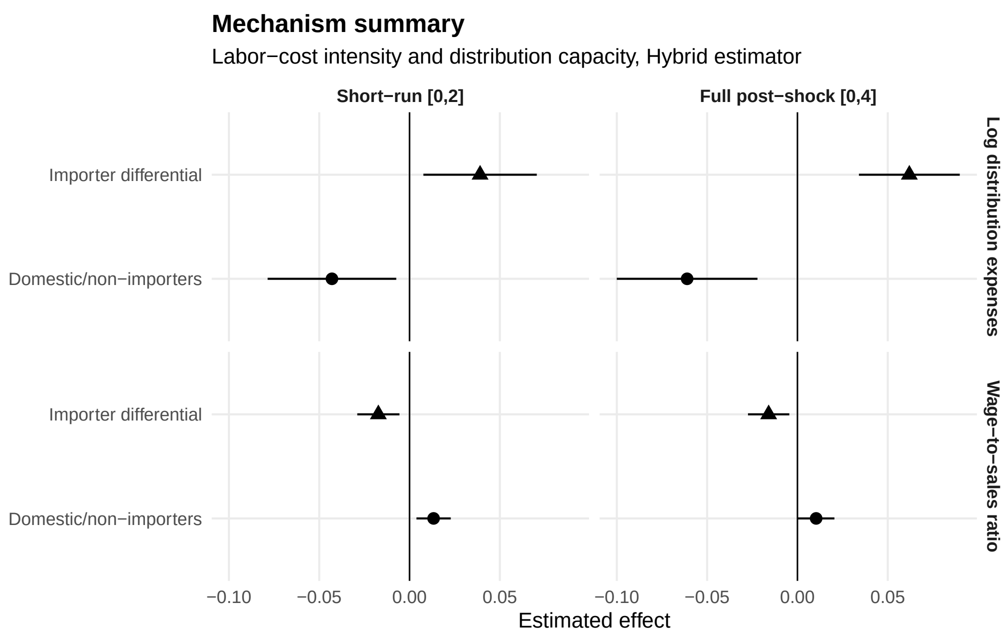

Figure 4: **Mechanism summary: labor-cost intensity and distribution capacity**

Notes: The figure summarizes the preferred Hybrid mechanism estimates for the post-treatment windows $[0, 2]$ and $[0, 4]$. The domestic/non-importer coefficients capture the effect of flood exposure among firms without pre-shock import links. The importer differential captures the additional effect for firms with pre-shock import links. Confidence intervals are based on standard errors clustered by firm and state.

This mechanism is closest to the supply-chain resilience literature. Todo, Nakajima, and Matous study the Great East Japan Earthquake and show that firm resilience after a natural disaster depends on the structure of supply-chain networks (Todo et al., 2015). My results point to a related margin in the context of Indian flood shocks: firms with pre-existing import links are better able to preserve profitability and distribution activity when local conditions deteriorate. This also complements work on shock propagation through production networks. Barrot and Sauvagnat show that supplier-specific input relationships can transmit idiosyncratic shocks to downstream firms (Barrot and Sauvagnat, 2016), while Boehm, Flaaen, and Pandalai-Nayar and Carvalho et al. show that the 2011 Tohoku earthquake propagated through input-output and supply-chain linkages (Boehm et al., 2019; Carvalho et al., 2021). The difference in my setting is that import links appear less as a channel of exposure and more as a form of flexibility: they help firms absorb a domestic flood shock by limiting cost pressure and preserving commercial capacity.

Finally, the importer result is consistent with the literature on imported intermediates and firm performance. Amiti and Konings show that access to imported intermediate inputs can raise firm productivity, while Halpern, Koren, and Szeidl emphasize the productivity gains from imported input varieties (Amiti and Konings, 2007; László Halpern, 2015). My results suggest that these import links may also matter during shocks: beyond improving average performance, they can provide firms with a margin of resilience when domestic sourcing or distribution conditions are disrupted.

The analysis above uses the preferred Hybrid estimator. Section 6 shows that the main cash-profit-margin pattern is robust to alternative estimators, richer controls, state-year fixed effects, alternative clustering choices, leave-one-event-out checks, winsorization choices, and placebo importer-status assignments.

6 Robustness and diagnostics

This section summarizes the robustness and diagnostic checks reported below and in the appendix. The goal is to assess whether the main cash-profit-margin result is driven by a particular estimator, comparison group, outcome construction, event, or inference choice.

6.1 Robustness summary

Table 6 shows that the main cash-profit-margin pattern is stable across specifications. Across alternative estimators, richer controls, and state-year fixed effects, the domestic/non-importer coefficient remains negative, while the importer differential remains positive and similar in magnitude. The same pattern holds in leave-one-event-out checks and under alternative winsorization thresholds. This suggests that the main profitability result is not driven by the estimator, a single flood episode, or the treatment of outliers.

Table 6: Robustness summary for cash-profit margins

	Post-treatment window [0,2]		Post-treatment window [0,4]	
Specification	Domestic/non-importers	Importer differential	Domestic/non-importers	Importer differential
Baseline Hybrid	-0.033***	0.035***	-0.027***	0.027***
	(0.010)	(0.008)	(0.010)	(0.008)
PS matching	-0.030***	0.030***	-0.026***	0.024***
	(0.006)	(0.006)	(0.007)	(0.006)
IPW	-0.026***	0.027***	-0.025***	0.021***
	(0.006)	(0.006)	(0.006)	(0.006)
Rich controls	-0.035***	0.032***	-0.029***	0.023***
	(0.011)	(0.008)	(0.008)	(0.008)
State-year fixed effects	-0.034***	0.033***	-0.027***	0.025***
	(0.009)	(0.008)	(0.009)	(0.008)
Sensitivity ranges				
Leave-one-event-out	$[-0.048, -0.028]$	$[0.029, 0.048]$	$[-0.037, -0.023]$	$[0.023, 0.037]$
Winsorization thresholds	$[-0.036, -0.024]$	$[0.025, 0.039]$	$[-0.028, -0.021]$	$[0.021, 0.030]$

Notes: The table summarizes robustness checks for the cash-profit-margin outcome. The domestic/non-importer coefficient is the average post-flood effect for firms without pre-shock import links. The importer differential is the additional effect for pre-shock importers. Standard errors clustered by firm and state are reported in parentheses for exact specifications. Leave-one-event-out and winsorization rows report the range of coefficients across robustness variants. Full results, including implied importer effects, observations, fixed effects, and additional outcomes, are reported in the appendix. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Across alternative estimators, richer controls, and state-year fixed effects, the domestic/non-importer coefficient remains negative while the importer differential remains positive and similar in magnitude. The same pattern holds in leave-one-event-out checks and under alternative winsorization thresholds. These results indicate that the main profitability pattern is not driven by a particular estimator, a single flood episode, or the treatment of outliers.

Appendix Tables 11 and 12 report the full outcome decomposition using propensity-score matching and IPW. The cash-profit-margin result is the most stable across estimators. The results for IHS-transformed cash profits are directionally similar but less robust, while effects on log sales and log raw-material use are smaller and less precisely estimated. This supports the interpretation that importer resilience is more visible in margins than in output scale.

6.2 Identification diagnostics

Appendix Table 13 reports joint pre-trend tests for all main outcomes and estimators. For the main outcome, cash-profit margin, the tests do not reject the null of no differential pre-trends for the baseline non-importer effects or the importer-differential effects across the Hybrid, propensity-score matching, and IPW specifications. The evidence is weaker for IHS-transformed cash profits, which is why the analysis places greater weight on cash-profit margins than on profit levels.

Table 7 reports additional pre-trend diagnostics for the preferred Hybrid specification. The close-lead tests over $[-3, -2]$ do not reject the null for the baseline, importer-differential, or implied-importer effects. Over the longer lead window $[-5, -2]$, however, the implied-importer test rejects the null at conventional levels. I therefore interpret the longer-lead evidence as a diagnostic caveat and focus the main event-study figure on the closer pre-treatment window, where event support is more comparable.

Table 7: Pre-trend diagnostics for cash-profit margins

Lead window	Baseline	Importer differential	Implied importer
$[-5, -2]$	0.627	0.341	0.028
$[-3, -2]$	0.628	0.887	0.282

Notes: The table reports joint Wald-test p-values for pre-treatment coefficients in the preferred Hybrid specification. Baseline refers to non-importer treatment leads. Importer differential refers to the additional pre-treatment leads for pre-shock importers. Implied importer refers to the sum of the baseline and importer-differential leads.

This distinction matters because the far leads are based on less comparable event support across flood episodes, whereas the close-lead window corresponds more directly to the pre-treatment periods used to assess the main event-study dynamics.

Appendix Tables 14 and 15 report the full rich-control and state-year fixed-effect specifications. These checks address the concern that flood exposure may be correlated with time-varying state-level shocks, local economic conditions, or observable pre-shock firm characteristics.

6.3 Timing, event dependence, and inference

Because flood exposure is an event-specific shock rather than an absorbing treatment, the standard distinction between never-treated and not-yet-treated controls in staggered-adoption designs does not map directly onto this setting. In the baseline stacked design, control firms are firms not exposed to flood event e within the corresponding event stack. Appendix Table 16 separates short-run effects over $[0, 2]$ from later post-shock effects over $[3, 4]$, while Appendix Table 17 shows that the main result is not driven by any single flood episode. This is important because the baseline analysis pools four major flood events.

Appendix Table 18 reports alternative clustering choices. The main conclusion remains similar when standard errors are clustered at the firm level, state level, or two-way firm-by-state level. This indicates that the statistical significance of the main result is not an artifact of a particular inference correction.

6.4 Placebo, heterogeneity, and mechanism diagnostics

Appendix Table 20 reports a placebo test based on fake importer status. The actual short-run importer differential is far from the placebo distribution: the actual estimate is 0.035, compared with a placebo mean of 0.002, and the randomization p-value is below 0.01. This weakens the concern that the importer differential is mechanically generated by the estimation design rather than by meaningful pre-shock differences in firms' sourcing capacity.

Finally, Appendix Table 21 reports sector-specific estimates. The importer-resilience pattern is not confined to a single sector, although precision varies across sector groups. Appendix Tables 22 and 23 report additional mechanism diagnostics, which support the interpretation that importers preserve margins through lower labor-cost pressure and stronger distribution or commercial capacity after flood exposure.

7 Discussion

The results point to a clear pattern: flood exposure reduces profitability among firms without pre-shock import links, while pre-shock importers are largely insulated from this decline. The effect is strongest for cash-profit margins. This is important because the resilience margin does not appear primarily through higher sales or larger input use after floods. Instead, the evidence suggests that importers are better able to preserve margins when local conditions deteriorate.

This interpretation matters for how we think about firm-level exposure to climate and infrastructure shocks. A natural reading of flood exposure is that it directly disrupts production by damaging local infrastructure, interrupting input availability, or reducing access to markets. The results are consistent with this view for non-importing firms. However, the importer differential suggests that exposure to the same local shock does not translate into the same profitability loss for all firms. Pre-existing access to international sourcing relationships appears to provide a form of operational flexibility. Importers may be better able to substitute away from disrupted local inputs, rely on more diversified supplier relationships, or maintain commercial activity when domestic supply conditions worsen.

The mechanism results support this interpretation. Flood exposure increases the wage-to-sales ratio of non-importers and reduces their distribution-related activity. By contrast, the importer differential moves in the opposite direction, offsetting much of these changes. This suggests that non-importers experience both margin compression and a contraction in operating capacity, while importers are better able to preserve the cost and distribution margins that sustain profitability. The evidence therefore points less to a story of post-flood expansion and more to one of loss avoidance.

The findings also contribute to the literature on production networks and shock propagation. Existing work emphasizes that shocks can travel through supplier-customer relationships and input-output linkages (Barrot and Sauvagnat, 2016; Boehm et al., 2019; Carvalho et al., 2021). This paper highlights a complementary margin. In the context of localized flood shocks, import links do not only represent potential exposure to external disruptions; they can also provide resilience when domestic conditions deteriorate. This connects the disaster-shock literature to work on imported intermediates, which shows that imported inputs can improve firm performance by expanding input choice, quality, and productivity (Amiti and Konings, 2007; László Halpern, 2015). The results suggest that these import links may matter not only in normal times, but also during periods of environmental disruption.

Several limitations are important. First, importer status is not randomly assigned. Although the empirical strategy compares exposed and non-exposed firms within a matched stacked event-study framework and controls for pre-shock firm characteristics, importer status may still capture unobserved differences in managerial capacity, supplier relationships, or organizational resilience. The alternative-resilience checks reduce this concern by showing that importer status predicts resilience more strongly than exporter status, size, assets, TFP, leverage, or age in the short run, but they do not fully eliminate it. The results should therefore be interpreted as evidence that pre-shock import links are associated with resilience, rather than as a purely causal estimate of becoming an importer.

Second, the flood-exposure measure captures major flood episodes using state-level event information combined with localized rainfall anomalies. This improves spatial precision relative to purely state-level exposure, but it does not directly observe firm-level flood damage, plant-level inundation, or disruptions to specific roads, suppliers, or distribution routes. As a result, the estimates should be interpreted as the effect of exposure to severe local flood conditions rather than the effect of verified physical damage at each firm.

Third, the evidence is strongest for cash-profit margins. Results for IHS-transformed cash profits are directionally similar but less robust, and the effects on sales and raw-material use are weaker. This pattern is informative, because it suggests that resilience operates through margins rather than scale. At the same time, it means that the paper's strongest conclusion concerns profitability preservation, not broader firm growth or output expansion.

Finally, the external validity of the results depends on the institutional and economic context. India during the sample period is a setting in which firms differ substantially in their access to imported inputs, infrastructure, and distribution networks. The importer-resilience margin may be especially relevant in environments where domestic supply chains are vulnerable to localized climate shocks and where import access provides meaningful sourcing flexibility. The effect may be weaker in contexts where imports are themselves highly disrupted, where customs or logistics constraints limit substitution, or where firms depend on highly specific domestic inputs.

The policy implications are therefore subtle. The results do not imply that import dependence is always protective, nor that exposure to global markets is costless. Rather, they suggest that firms with more diversified sourcing options may be better able to absorb localized environmental shocks. Policies that improve firms' access to reliable input markets, logistics infrastructure, and diversified supplier networks may therefore increase resilience to climate-related disruptions. In this sense, adaptation is not only a matter of physical protection against floods; it also depends on the flexibility of the economic networks through which firms source inputs and reach markets.

8 Conclusion

This paper studies whether pre-shock import links help firms withstand local flood shocks. Using a stacked event-study design that combines firm-level panel data with major flood episodes in India, I estimate the effect of flood exposure on firm profitability and examine whether the response differs between importers and non-importers.

The results show that flood exposure reduces the profitability of non-importing firms. In the preferred Hybrid specification, the short-run cash-profit margin of non-importers falls by 3.3 percentage points. By contrast, the importer differential is positive and similar in magnitude, leaving the implied effect for pre-shock importers close to zero. The same pattern holds over the full post-shock window. The effect is strongest for cash-profit margins, directionally similar for cash-profit levels, and weaker for sales and raw-material use.

The mechanism evidence suggests that importer resilience operates through the preservation of margins and operating capacity. Non-importers experience higher labor-cost intensity and lower distribution-related activity after flood exposure. Importers, by contrast, appear less exposed to these forms of margin compression. These findings suggest that pre-existing import links provide firms with a margin of flexibility when local supply and distribution conditions are disrupted.

The paper contributes to the literature on natural disasters, production networks, and firm resilience by showing that supply-chain relationships can matter not only for the propagation of shocks, but also for firms' ability to absorb them. In the context of localized flood shocks, import links appear to operate as a form of resilience rather than simply as a source of exposure. More broadly, the results suggest that adaptation to climate-related shocks depends not only on physical infrastructure and local protection, but also on the structure and flexibility of firms' sourcing networks.

References

Abadie, A. and Imbens, G. W. (2006). Large sample properties of matching estimators for average treatment effects. *Econometrica*, 74(1):235–267.

Abadie, A. and Spiess, J. (2022). Robust post-matching inference. *Journal of the American Statistical Association*, 117(538):983–995.

Albala-Bertrand, J. (1993). Natural disaster situations and growth: A macroeconomic model for sudden disaster impacts. *World Development*, 21(9):1417–1434.

Amiti, M. and Konings, J. (2007). Trade liberalization, intermediate inputs, and productivity: Evidence from indonesia. *American Economic Review*, 97(5):1611–1638.

Austin, P. C. (2009). Balance diagnostics for comparing the distribution of baseline covariates between treatment groups in propensity-score matched samples. *Statistics in Medicine*, 28(25):3083–3107.

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.

Bas, M. and Paunov, C. (2024). Riders on the storm: How do firms navigate production and market conditions amid el nino? *Journal of Development Economics*, 172.

Bernard, A. B., Jensen, J. B., Redding, S. J., and Schott, P. K. (2007). Firms in international trade. *Journal of Economic Perspectives*, 21(3):105–130.

Bloom, N., Eifert, B., Mahajan, A., McKenzie, D., and Roberts, J. (2013). Does management matter? evidence from india. *Quarterly Journal of Economics*, 128(1):1–51.

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. *Review of Economics and Statistics*, 101(1):60–75.

Borusyak, K., Jaravel, X., and Spiess, J. (2024). Revisiting event-study designs: Robust and efficient estimation. *Review of Economic Studies*, 91(6):3253–3285.

Callaway, B. and Sant'Anna, P. H. C. (2021). Difference-in-differences with multiple time periods. *Journal of Econometrics*, 225(2):200–230.

Cameron, A. C., Gelbach, J. B., and Miller, D. L. (2011). Robust inference with multiway clustering. *Journal of Business & Economic Statistics*, 29(2):238–249.

Carvalho, V. M., Nirei, M., Saito, Y. U., and Tahbaz-Salehi, A. (2021). Supply chain disruptions: Evidence from the great east japan earthquake. *Quarterly Journal of Economics*, 136(2):1255–1321.

Cengiz, D., Dube, A., Lindner, A., and Zipperer, B. (2019). The effect of minimum wages on low-wage jobs. *Quarterly Journal of Economics*, 134(3):1405–1454.

de Mel, S., McKenzie, D., and Woodruff, C. (2012). Enterprise recovery following natural disasters. *The Economic Journal*, 122(559):64–91.

Diamond, A. and Sekhon, J. S. (2013). Genetic matching for estimating causal effects: A general multivariate matching method for achieving balance in observational studies. *The Review of Economics and Statistics*, 95(3):932–945.

Felix, K.T.; Balasubramanian, M. G. P. K. B. (2025). Assessing the socioeconomic and environmental determinants of flood vulnerability in india: a panel data approach. *Scientic Report*, 15.

Goldberg, P. K., Khandelwal, A. K., Pavcnik, N., and Topalova, P. (2010). Imported intermediate inputs and domestic product growth: Evidence from india. *Quarterly Journal of Economics*, 125(4):1727–1767.

Hirano, K., Imbens, G. W., and Ridder, G. (2003). Efficient estimation of average treatment effects using the estimated propensity score. *Econometrica*, 71(4):1161–1189.

Intergovernmental Panel on Climate Change (2021). Climate change 2021: The physical science basis (working group i). *Newsroom*.

Kahn, M. E. (2005). The death toll from natural disasters: The role of income, geography, and institutions. *The Review of Economics and Statistics*, 87:271–284.

Lai, S., Chen, L., Wang, Q. S., and Anderson, H. (2022). Natural disasters, trade credit, and firm performance. *Economic Modelling*, 116.

Loayza, N. V., Olaberría, E., Rigolini, J., and Christiaensen, L. (2012). Natural disasters and growth: Going beyond the averages. *World Development*, 40(7):1317–1336.

László Halpern, Miklós Koren, A. S. (2015). Imported inputs and productivity. *American Economic Review*, 105(12):3660–3703.

Manova, K. (2013). Credit constraints, heterogeneous firms, and international trade. *Review of Economic Studies*, 80(2):711–744.

Nedoncelle, C. and Wolfersberger, J. (2024). Weather shocks and firm exports in developing countries. *Review of World Economics*.

Noy, I. (2009). The macroeconomic consequences of disasters. *Journal of Development Economics*, 88(2):221–231.

Ritzwoller, D. M., Romano, J. P., and Shaikh, A. M. (2024). Randomization inference: Theory and applications. *Forthcoming in Journal of Political Economy Microeconomics*.

Rosenbaum, P. R. and Rubin, D. B. (1983). The central role of the propensity score in observational studies for causal effects. *Biometrika*, 10(1):41–55.

Roth, J. (2022). Pretest with caution: Event-study estimates after testing for parallel trends. *American Economic Review: Insights*, 4(3):305–322.

Stuart, E. A. (2010). Matching methods for causal inference: A review and a look forward. *Statistical Science*, 25(1):1–21.

Sun, L. and Abraham, S. (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics*, 225(2):175–199.

Todo, Y., Nakajima, K., and Matous, P. (2015). How do supply chain networks affect the resilience of firms to natural disasters? evidence from the great east japan earthquake. *Journal of Regional Science*, 55(2):209–229.

Wing, C., Freedman, S. M., and Hollingsworth, A. (2024). Stacked difference-in-differences. NBER Working Paper 32054, National Bureau of Economic Research.

9 Appendix

This appendix provides supporting material for the empirical analysis. I first document the archival and rainfall sources used to construct flood exposure. I then report balance diagnostics and estimator-specific sample support for the matched and weighted samples. The remaining subsections provide the full robustness tables referenced in the main text, including alternative estimators, pre-trend tests, richer controls, state-year fixed effects, timing checks, alternative clustering, leave-one-event-out estimates, winsorization checks, placebo tests, sector heterogeneity, and mechanism diagnostics.

9.1 Exposure documentation and archival sources

This subsection documents the archival and rainfall sources used to construct flood exposure. The Indian Meteorological Department's *Disastrous Weather Events* reports provide narrative descriptions of major flood episodes, including their timing, affected regions, and reported damages. I use these reports to identify the affected states $\mathcal{S}_e$ and to select the peak flood month $m(e)$ for each event-year, cross-checking the event classification against EM-DAT. The archival information therefore provides an administrative and meteorological anchor for the subsequent rainfall-based exposure assignment.

Figures 5 and 6 reproduce the presentation page and an illustrative extract from the 1993 report. They show the type of archival information used to identify the timing and geographic extent of the flood episode.

Figure 5: Title page of the 1993 *Disastrous Weather Events* report

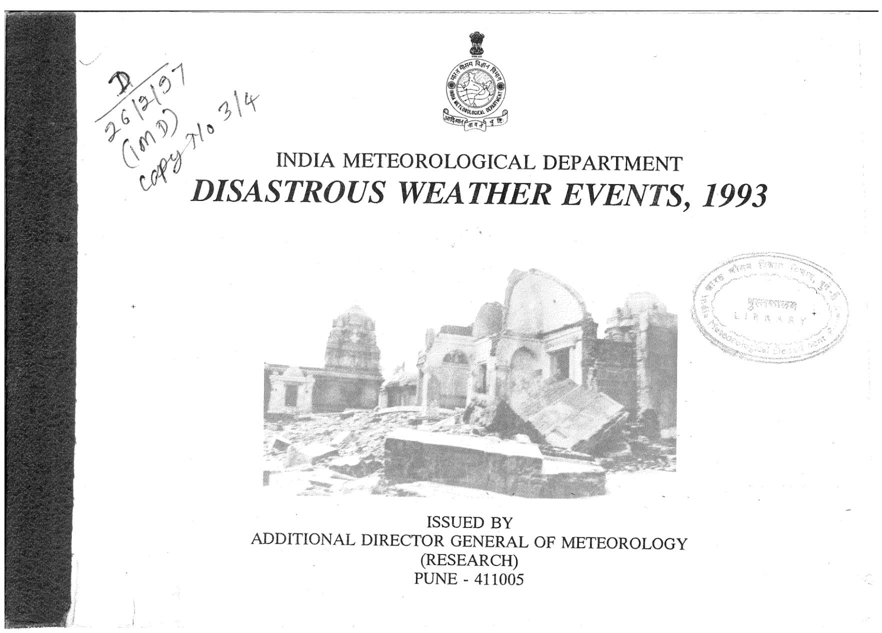

INDIA METEOROLOGICAL DEPARTMENT

DISASTROUS WEATHER EVENTS, 1993

ISSUED BY
ADDITIONAL DIRECTOR GENERAL OF METEOROLOGY
(RESEARCH)
PUNE - 411005

Figure 6: Illustrative flood record from the 1993 *Disastrous Weather Events* report

S.N.	Date/period	Area affected	Intensity (river)	Casualties	Extent of damage
			TABLE - XI : FLOODS AND HEAVY RAINS		
		STATES			
		ANDHRA PRADESH			
1.	7-11 Oct.	Hyderabad, Kurnool and Mehboobnagar	Flash flood	13 persons died in Hyderabad.	(i) 2,265 houses fully and 8,149 partially damaged. (ii) 109 villages affected. (iii) Extensive damage of standing crop reported.
		ARUNACHAL PRADESH			
1.	7 Jun.	Entire State	Heavy rains		Vehicular traffic disrupted due to landslide at several places.
2.	2-7 Jul.	Lohit and Lower Subansiri	Severe flood (Noa Dihing)		(i) 4 houses washed away. (ii) 4 bridges on Yazadazero road washed away. (iii) Both the districts cut off from the rest of the country. (iv) 150 metre stretch of national highway near Krishnapur in Namsai circle flooded. (v) National highway No. 52 submerged at several places.
3.	5 Sept.	Tawang	Flash flood	14 persons died and one injured.	10 houses washed away.

19

The archival sources identify the states affected by each event, but state-level classifications remain too coarse to assign exposure at the firm level. I therefore combine the archival classification with spatial variation in rainfall intensity. Figure 7 presents the standardized CHIRPS precipitation anomalies for the remaining flood events analyzed in the paper. Each map measures precipitation in the event-specific peak month relative to the pixel-specific long-run distribution for the same calendar month over 1981–2025. Red borders indicate states classified as affected using the IMD archives and EM-DAT.

The maps show substantial within-state variation in rainfall intensity. Some areas inside affected states experienced large positive anomalies, while others received rainfall much closer to their historical distribution. This illustrates why assigning treatment to every firm within an affected state would introduce substantial exposure misclassification. The treatment construction therefore classifies firms as exposed only when they are located in pixels in the upper tail of the rainfall-anomaly distribution within an affected state, as described in Section 3.2.

Figure 7: Standardized precipitation anomalies for the flood events

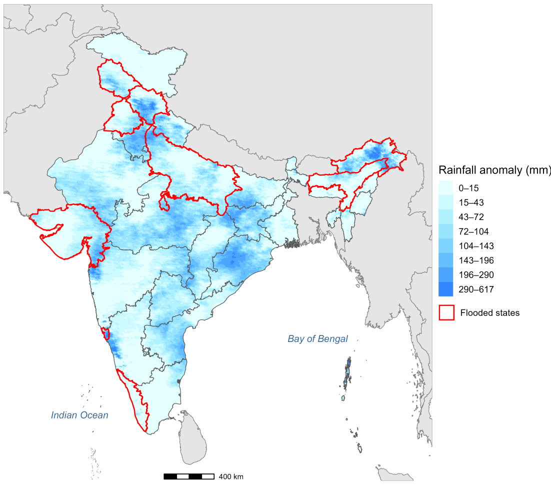

(a) August 1994

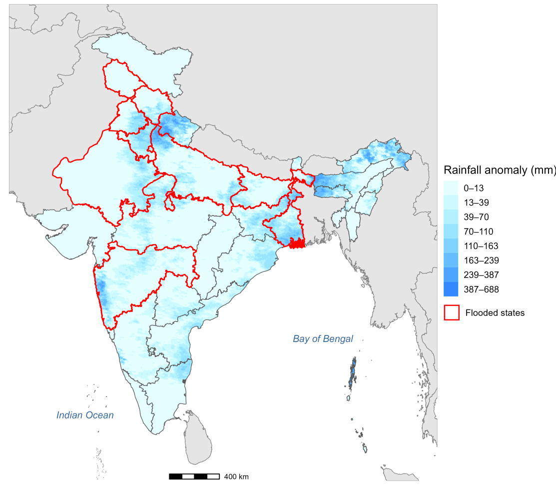

(b) September 1995

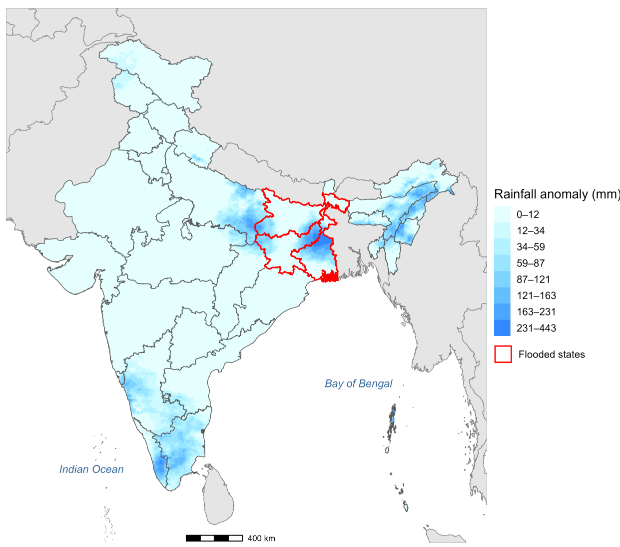

(c) September 2000

Taken together, the archival reports and rainfall anomaly maps document the two-step exposure procedure: flood episodes are first anchored in historically documented affected states, after which firm exposure is assigned using localized rainfall anomalies within those states.

9.2 Balance diagnostics and estimator support

This subsection evaluates whether the matching and weighting procedures improve the comparability of flood-exposed and non-exposed firms in terms of observed pre-treatment characteristics. Propensity-score methods seek to construct treated and comparison groups with similar distributions of observed covariates (?). I assess covariate balance using absolute standardized mean differences (SMDs), which measure the magnitude of imbalance independently of sample size and are therefore preferable to conventional significance tests for this purpose (Austin, 2009).

Figures 8 and 9 report covariate-level balance for IPW and propensity-score matching, respectively. Table 8 summarizes balance across a broader set of pre-treatment outcomes and mechanism variables. Tables 9 and 10 then document the sample support and regression sample associated with each estimator.

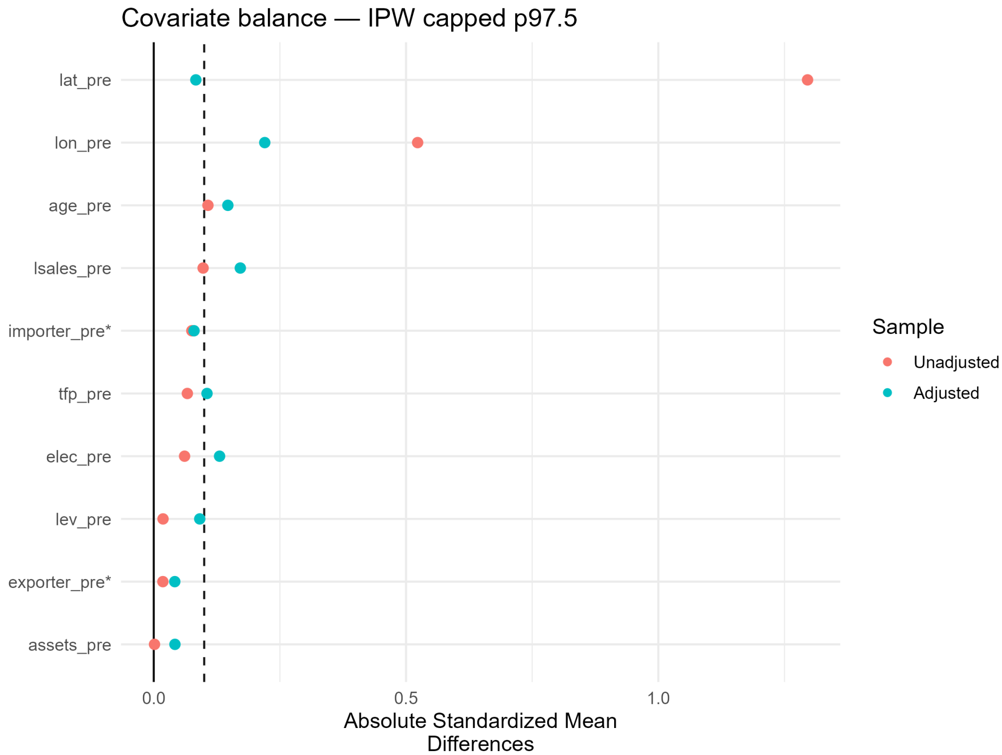

Figure 8: **Covariate balance for the IPW estimator**

Notes: The figure reports absolute standardized mean differences between flood-exposed and non-exposed firms before and after applying inverse probability weighting. Balance is assessed using pre-treatment covariates measured over the pre-shock window: log sales, assets, TFP, leverage, electricity use, firm age, importer status, exporter status, latitude, and longitude. The dashed vertical line marks the conventional 0.10 threshold for absolute standardized mean differences. Asterisks identify binary covariates, as reported by the balance-plot routine.

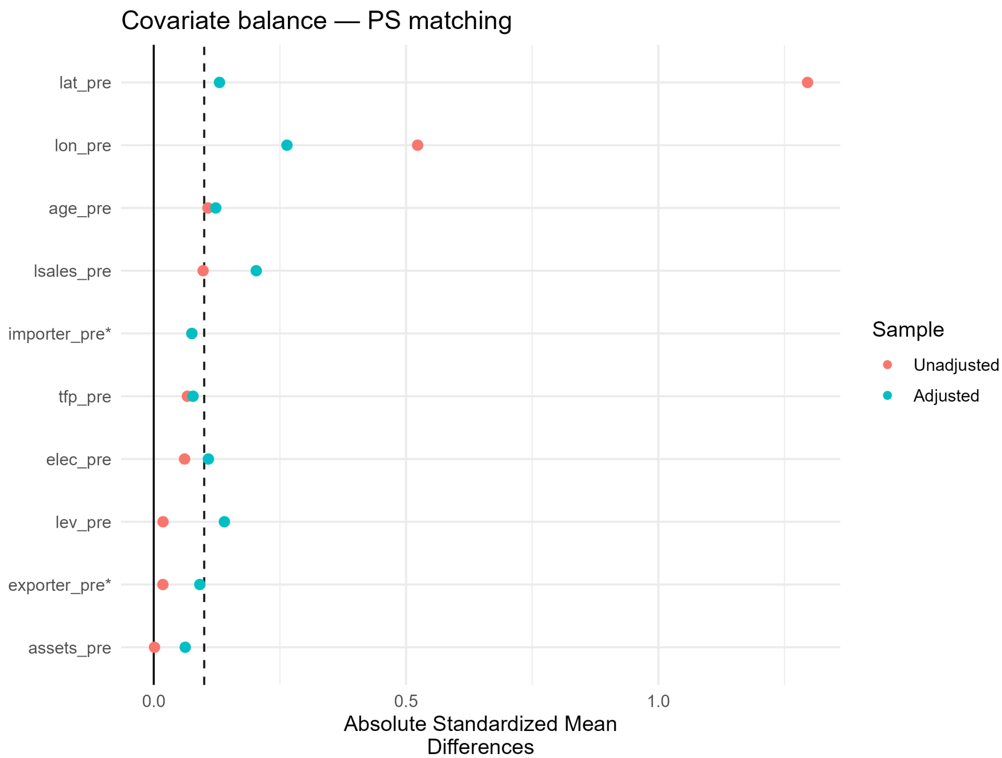

Figure 9: **Covariate balance for the PS matching estimator**

Notes: The figure reports absolute standardized mean differences between flood-exposed and non-exposed firms before and after propensity-score matching. Balance is assessed using pre-treatment covariates measured over the pre-shock window: log sales, assets, TFP, leverage, electricity use, firm age, importer status, exporter status, latitude, and longitude. The dashed vertical line marks the conventional 0.10 threshold for absolute standardized mean differences. Asterisks identify binary covariates, as reported by the balance-plot routine.

Table 8: Extended pre-treatment balance across estimators

Estimator	Covariates	Mean SMD before	Mean SMD after	Max SMD after	Share below 0.10 after
Hybrid	20	0.113	0.084	0.242	0.700
PS matching	20	0.113	0.138	0.338	0.450
IPW	20	0.108	0.101	0.257	0.550

Notes: The table reports extended balance diagnostics for pre-treatment outcomes and mechanism variables. SMD denotes the absolute standardized mean difference between exposed and non-exposed firms. The "before" column reports balance before matching or weighting, while the "after" columns report balance after applying the corresponding estimator-specific weights. The extended balance variables include pre-treatment sales, cash profits, cash-profit margins, raw-material use, input-cost ratios, marketing and distribution expenses, borrowing, and international exposure measures. These variables are used as diagnostic checks and do not redefine the frozen matching or weighting procedure. Hybrid refers to Mahalanobis matching with a propensity-score caliper; PS refers to propensity-score matching; IPW refers to inverse probability weighting with weights capped at the 97.5th percentile.

The extended diagnostics support the choice of the Hybrid estimator as the preferred specification. It produces the lowest mean post-adjustment SMD, at 0.084, and the largest share of covariates below the 0.10 threshold, at 70%. IPW also improves average balance relative to the unweighted sample, although its maximum remaining imbalance is larger and only 55% of the extended covariates fall below the threshold. By contrast, PS matching improves balance for some covariates but worsens the average SMD across the broader diagnostic set. These results motivate using the Hybrid estimator for the main analysis while retaining PS matching and IPW as alternative specifications.

Balance must be considered jointly with sample support. Matching can improve comparability by discarding observations outside the region of common support, whereas weighting retains a much larger fraction of the original sample. Table 9 therefore reports the number of stacked firm-year observations, firms, states, and treated and control observations retained by each estimator.

Table 9: Estimator-specific sample support

Estimator	Observations	Firms	States	Events	Treated obs.	Control obs.
Hybrid	11,518	963	24	4	8,471	3,047
PS	12,632	1,027	25	4	9,439	3,193
IPW	63,258	3,668	26	4	9,439	53,819

Notes: The table reports estimator-specific sample support for the cash-profit-margin outcome. The unit of observation is a stacked firm-year. Hybrid refers to the preferred Mahalanobis matching procedure with a propensity-score caliper. PS refers to propensity-score matching, and IPW refers to inverse probability weighting. Treated observations correspond to firm-years located in flood-exposed pixels within the relevant event stack; control observations correspond to non-exposed firm-years.

The Hybrid estimator retains 963 firms and 11,518 stacked firm-year observations. Its smaller sample reflects the stricter combination of Mahalanobis-distance matching and a propensity-score caliper. PS matching retains a somewhat larger sample, while IPW uses the broadest support because it reweights rather than discards most comparison observations. The larger IPW sample increases coverage, but it comes with weaker balance than the preferred Hybrid specification.

Table 10 reports the observations that enter the final cash-profit-margin regressions after accounting for outcome availability and fixed-effect estimation restrictions. It also reports overall and within-R^2 statistics for the corresponding specifications.

Table 10: Main model sample sizes and fit

Estimator	Model observations	R^2	Within R^2
Hybrid	11,479	0.414	0.004
PS	12,603	0.434	0.004
IPW	63,116	0.445	0.003

Notes: The table reports the number of observations and model-fit statistics for the main cash-profit-margin specifications. All specifications include firm, year, and event fixed effects. Differences between the support statistics in Table 9 and the model observations reported here reflect outcome availability and fixed-effect estimation restrictions.

The difference between estimator support and final model observations is small for all three estimators. The low within-R^2 values are not unusual in firm fixed-effect specifications, because they measure the share of within-firm variation explained after absorbing firm, year, and event fixed effects. The purpose of comparing the estimators is therefore not their relative model fit, but the trade-off between covariate balance and sample coverage.

9.3 Alternative estimators and pre-treatment diagnostics

This subsection examines whether the main results are sensitive to the method used to construct the comparison group. The preferred Hybrid estimator combines Mahalanobis-distance matching with a propensity-score caliper, thereby prioritizing covariate balance and common support. Propensity-score matching and inverse probability weighting provide two alternative approaches. PS matching constructs a matched comparison sample using estimated treatment probabilities, while IPW retains a substantially broader sample and reweights observations to make exposed and non-exposed firms comparable in observed pre-treatment characteristics. Comparing the results across these estimators therefore helps determine whether the main findings depend on a particular matching rule, weighting scheme, or estimation sample.

Tables 11 and 12 report the full outcome decomposition using propensity-score matching and IPW, respectively. The specifications otherwise follow the baseline event-study design and include firm, year, and event fixed effects.

Table 11: Main outcome results using PS matching

	Cash-profit margin	Asinh cash profits	Log sales	Log raw materials
Post-treatment window [0,2]				
Domestic/non-importers	-0.030***	-0.173***	-0.035	-0.033
	(0.006)	(0.026)	(0.023)	(0.033)
Importer differential	0.030***	0.174***	0.052**	0.058**
	(0.006)	(0.028)	(0.021)	(0.023)
Implied importer effect	0.000	0.002	0.017	0.025
Post-treatment window [0,4]				
Domestic/non-importers	-0.026***	-0.116***	-0.019	-0.018
	(0.007)	(0.028)	(0.031)	(0.051)
Importer differential	0.024***	0.077	0.029	0.042
	(0.006)	(0.050)	(0.026)	(0.036)
Implied importer effect	-0.003	-0.039	0.010	0.024
Observations	12,603	12,603	12,603	12,603
R^2	0.434	0.576	0.919	0.893
Within R^2	0.004	0.003	0.004	0.003
Firm fixed effects	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes
Event fixed effects	Yes	Yes	Yes	Yes

Notes: The table reports post-treatment average effects for the main outcomes using propensity-score matching (PS). Standard errors, reported in parentheses, are clustered by firm and state. Stars denote significance at the 10%, 5%, and 1% levels.

Table 12: Main outcome results using IPW

	Cash-profit margin	Asinh cash profits	Log sales	Log raw materials
Post-treatment window [0,2]				
Domestic/non-importers	-0.026***	-0.144***	-0.032	-0.031
	(0.006)	(0.022)	(0.021)	(0.030)
Importer differential	0.027***	0.155***	0.048**	0.053*
	(0.006)	(0.028)	(0.019)	(0.030)
Implied importer effect	0.001	0.012	0.015	0.023
Post-treatment window [0,4]				
Domestic/non-importers	-0.025***	-0.102***	-0.016	-0.022
	(0.006)	(0.034)	(0.025)	(0.043)
Importer differential	0.021***	0.057	0.024	0.038
	(0.006)	(0.042)	(0.024)	(0.041)
Implied importer effect	-0.004	-0.045	0.008	0.016
Observations	63,116	63,116	63,116	63,116
R^2	0.445	0.586	0.917	0.893
Within R^2	0.003	0.002	0.003	0.002
Firm fixed effects	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes
Event fixed effects	Yes	Yes	Yes	Yes

Notes: The table reports post-treatment average effects for the main outcomes using inverse probability weighting (IPW). Standard errors, reported in parentheses, are clustered by firm and state. Stars denote significance at the 10%, 5%, and 1% levels.

The cash-profit-margin result is highly stable across the alternative estimators. Under PS matching, the short-run domestic/non-importer effect is -0.030, while the importer differential is 0.030. Under IPW, the corresponding estimates are -0.026 and 0.027. The same offsetting pattern remains visible over the full $[0, 4]$ window, and the implied importer effect remains close to zero in both specifications. These estimates are similar in sign and magnitude to the preferred Hybrid results, despite meaningful differences in balance and sample support across estimators.

The results for IHS-transformed cash profits are also similar in the short run: non-importers experience a decline, while the importer differential offsets most of the loss. Over the full post-treatment window, however, the importer differential becomes smaller and less precisely estimated. The sales and raw-material results are weaker and generally imprecise. Taken together, the alternative estimators reinforce the conclusion that the most robust resilience margin concerns profitability rather than post-flood expansion in sales or input use.

The alternative-estimator results are informative only if exposed and comparison firms do not display systematically different outcome dynamics before the flood. Table 13 therefore reports joint Wald tests of the pre-treatment event-study coefficients over the lead window $[-5, -2]$. The baseline test examines whether the pre-flood coefficients for domestic non-importers are jointly equal to zero. The differential test examines whether the additional pre-flood coefficients for importers are jointly equal to zero.

These tests provide a diagnostic for systematic differences in pre-treatment dynamics, but failure to reject the null should not be interpreted as proof of parallel trends. Pre-trend tests may have limited power against economically relevant violations, and their interpretation should therefore be combined with the graphical event-study evidence and the stability of the estimates across specifications (Roth, 2022).

Table 13: Pre-trend tests for main outcomes, all estimators

	Hybrid		PS		IPW	
Outcome	Baseline p	Diff. p	Baseline p	Diff. p	Baseline p	Diff. p
Cash-profit margin	0.627	0.341	0.624	0.200	0.700	0.141
Asinh cash profits	0.098	0.031	0.051	0.041	0.415	0.114
Log sales	0.650	0.230	0.584	0.176	0.482	0.108
Log raw materials	0.560	0.151	0.466	0.113	0.367	0.135

Notes: The table reports joint Wald-test p-values for pre-treatment event-study coefficients over the lead window $[-5, -2]$. Baseline refers to the coefficients for domestic non-importers, while Diff. refers to the additional coefficients for pre-shock importers. The omitted reference period is event time $k = -1$. The table reports separate tests for the baseline and importer-differential coefficients; tests of the implied importer effect are reported in Table 7.

The cash-profit-margin outcome displays the strongest pre-treatment support. Across the Hybrid, PS, and IPW specifications, the tests do not reject the null for either the domestic/non-importer leads or the importer-differential leads. The same broad pattern holds for log sales and log raw-material use, for which none of the joint tests rejects at conventional significance levels.

The evidence is weaker for IHS-transformed cash profits. In the Hybrid and PS specifications, the importer-differential tests reject the null at the 5% level, while the baseline tests are either marginal or close to conventional significance thresholds. The corresponding IPW tests do not reject, but the variation across estimators warrants caution. I therefore treat the cash-profit-level estimates as supportive evidence and place greater weight on the cash-profit-margin outcome, for which both the coefficient pattern and the pre-treatment diagnostics are consistently stronger.

9.4 Additional robustness checks

This subsection evaluates the sensitivity of the preferred Hybrid estimates to several potential threats. The checks address four distinct concerns: omitted firm- and state-level factors, the timing and event composition of the estimated effects, the treatment of statistical inference and outcome outliers, and the possibility that the importer differential is mechanically generated by the estimation procedure. Each test therefore targets a different aspect of the empirical design rather than simply re-estimating the baseline model under arbitrary alternatives.

9.4.1 Additional controls and state-year shocks

The first set of checks examines whether the estimated importer-resilience margin reflects observed differences in pre-shock firm characteristics or contemporaneous shocks at the state level. Table 14 allows treatment dynamics to vary with a richer set of pre-shock firm characteristics. This specification addresses the concern that flood-exposed importers and non-importers may differ systematically along dimensions that also shape their post-flood performance. If the main coefficients changed substantially after introducing these interactions, the baseline importer differential could be capturing broader differences in firm composition rather than resilience associated with import links.

Table 14: Robustness to richer controls, Hybrid estimator

	Cash-profit margin	Asinh cash profits	Log sales	Log raw materials
Post-treatment window [0,2]				
Domestic/non-importers	-0.035***	-0.253***	-0.056	-0.012
	(0.011)	(0.065)	(0.041)	(0.068)
Importer differential	0.032***	0.213***	0.027	0.024
	(0.008)	(0.073)	(0.022)	(0.031)
Implied importer effect	-0.003	-0.039	-0.029	0.013
Post-treatment window [0,4]				
Domestic/non-importers	-0.029***	-0.251***	-0.039	-0.002
	(0.008)	(0.055)	(0.057)	(0.091)
Importer differential	0.023***	0.110	0.006	0.016
	(0.008)	(0.109)	(0.028)	(0.038)
Implied importer effect	-0.006	-0.141	-0.034	0.014
Observations	11,479	11,479	11,479	11,479
R^2	0.420	0.587	0.914	0.887
Within R^2	0.013	0.034	0.023	0.019
Firm fixed effects	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes
Event fixed effects	Yes	Yes	Yes	Yes

Notes: The specification allows treatment dynamics to vary with pre-shock firm characteristics, including firm size, assets, productivity, leverage, age, and international exposure. Standard errors are clustered by firm and state. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

The cash-profit-margin estimates remain close to the baseline coefficients. The domestic/non-importer effect remains negative and statistically significant, while the importer differential remains positive and significant in both post-treatment windows. This stability reduces the concern that the main result reflects differential responses associated with the observed pre-shock firm characteristics included in the specification.

Table 15 instead addresses time-varying shocks common to firms located in the same state and year. State-year fixed effects absorb state-level economic conditions, policies, infrastructure disruptions, and other contemporaneous shocks that could be correlated with both flood exposure and firm outcomes. Identification in this specification therefore comes from differences in localized flood exposure among firms observed within the same state-year.

Table 15: Robustness to state-year fixed effects, Hybrid estimator

	Cash-profit margin	Asinh cash profits	Log sales	Log raw materials
Post-treatment window [0,2]				
Domestic/non-importers	-0.034***	-0.166***	0.004	0.002
	(0.009)	(0.053)	(0.026)	(0.039)
Importer differential	0.033***	0.188***	0.031	0.034
	(0.008)	(0.037)	(0.020)	(0.029)
Implied importer effect	-0.001	0.021	0.035	0.036
Post-treatment window [0,4]				
Domestic/non-importers	-0.027***	-0.074	0.036	0.038
	(0.009)	(0.077)	(0.040)	(0.059)
Importer differential	0.025***	0.081	0.005	0.014
	(0.008)	(0.062)	(0.028)	(0.044)
Implied importer effect	-0.003	0.007	0.041	0.052
Observations	11,459	11,459	11,459	11,459
R^2	0.447	0.592	0.916	0.890
Within R^2	0.004	0.003	0.004	0.004
Firm fixed effects	Yes	Yes	Yes	Yes
State-Year fixed effects	Yes	Yes	Yes	Yes
Event fixed effects	Yes	Yes	Yes	Yes

*Notes :*The specification includes firm, event, and state-year fixed effects. Standard errors are clustered by firm and state. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

The cash-profit-margin coefficients remain similar to the baseline estimates. This reduces the concern that the main result is driven by unobserved state-level conditions occurring in the same years as the flood events.

9.4.2 Timing and event dependence

The next checks examine whether the estimated effects are concentrated immediately after the flood and whether the pooled result is disproportionately driven by one of the four events. Table 16 separates the immediate post-flood period $[0, 2]$ from the later period $[3, 4]$. This distinction helps determine whether the estimates represent a temporary disruption, a persistent profitability effect, or an average that masks different short- and medium-run responses.

Table 16: Persistence check: short-run and later post-flood windows, Hybrid estimator

	Cash-profit margin	Asinh cash profits	Log sales	Log raw materials
Post-treatment window [0,2]				
Domestic/non-importers	-0.033***	-0.196***	-0.031	-0.023
	(0.010)	(0.044)	(0.029)	(0.040)
Importer differential	0.035***	0.213***	0.049**	0.052
	(0.008)	(0.035)	(0.021)	(0.032)
Implied importer effect	0.002	0.018	0.018	0.029
Post-treatment window [3,4]				
Domestic/non-importers	-0.018	-0.012	0.026	0.035
	(0.014)	(0.110)	(0.058)	(0.090)
Importer differential	0.015	-0.062	-0.013	0.003
	(0.013)	(0.130)	(0.043)	(0.071)
Implied importer effect	-0.003	-0.074	0.013	0.038
Post-treatment window [0,4]				
Domestic/non-importers	-0.027***	-0.122**	-0.008	0.000
	(0.010)	(0.059)	(0.039)	(0.059)
Importer differential	0.027***	0.103*	0.024	0.032
	(0.008)	(0.056)	(0.028)	(0.045)
Implied importer effect	0.000	-0.019	0.016	0.032
Observations	11,479	11,479	11,479	11,479
R^2	0.414	0.573	0.913	0.885
Within R^2	0.004	0.003	0.004	0.004
Firm fixed effects	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes
Event fixed effects	Yes	Yes	Yes	Yes

Notes: The specification include firm, year, and event fixed effects unless otherwise indicated. Standard errors are clustered by firm and state. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. The table separates the immediate post-flood window $[0, 2]$, the later window $[3, 4]$, and the full post-flood window $[0, 4]$.

The cash-profit-margin effect is concentrated in the event year and the following two years. In the later window $[3, 4]$, both the domestic/non-importer effect and the importer differential are smaller and imprecisely estimated. The significant estimates over the full $[0, 4]$ window are therefore primarily driven by the short-run response rather than by a separately significant effect in years three and four. This pattern suggests that the importer-resilience margin is strongest during the immediate aftermath of the flood and subsequently weakens.

Table 17 re-estimates the preferred specification four times, dropping one flood episode in each column. This exercise tests whether the pooled result reflects a general pattern across events or is driven by an unusually large or influential episode. Stability in the sign and approximate magnitude of the coefficients across columns would indicate that the headline result does not depend on any single flood event.

Table 17: Leave-one-event-out robustness for cash-profit margins

	1993	1994	1995	2000
Post-treatment window [0,2]				
Domestic/non-importers	-0.040***	-0.032***	-0.048***	-0.028**
	(0.011)	(0.009)	(0.018)	(0.012)
Importer differential	0.041***	0.034***	0.048***	0.029***
	(0.007)	(0.007)	(0.017)	(0.010)
Implied importer effect	0.001	0.002	0.000	0.001
Post-treatment window [0,4]				
Domestic/non-importers	-0.030**	-0.033***	-0.037**	-0.023*
	(0.012)	(0.010)	(0.017)	(0.012)
Importer differential	0.032***	0.028***	0.037**	0.023**
	(0.008)	(0.008)	(0.014)	(0.011)
Implied importer effect	0.002	-0.004	0.000	0.000
Firm fixed effects	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes
Event fixed effects	Yes	Yes	Yes	Yes

Notes: The specification include firm, year, and event fixed effects unless otherwise indicated. Standard errors are clustered by firm and state. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.Each column re-estimates the preferred Hybrid specification after excluding the flood event indicated in the column heading. The outcome is the cash-profit margin. Standard errors are clustered by firm and state.

The domestic/non-importer coefficient remains negative and the importer differential remains positive in every leave-one-event-out specification. The implied importer effect also remains close to zero. The pooled result is therefore not attributable to one particular flood episode.

9.4.3 Inference and outcome construction

Table 18 examines whether statistical inference is sensitive to the assumed structure of residual dependence. Firm-level clustering allows errors to be correlated over time within firms, state-level clustering allows common dependence among firms located in the same state, and two-way clustering accommodates both dimensions simultaneously (Cameron et al., 2011). Because the clustering choice affects the estimated variance rather than the point estimates, this exercise tests the robustness of statistical significance rather than coefficient stability.

Table 18: Alternative clustering robustness for cash-profit margins

	Firm	State	Firm + state
Post-treatment window [0,2]			
Domestic/non-importers	-0.033***	-0.033***	-0.033***
	(0.010)	(0.010)	(0.010)
Importer differential	0.035***	0.035***	0.035***
	(0.011)	(0.008)	(0.008)
Implied importer effect	0.002	0.002	0.002
Post-treatment window [0,4]			
Domestic/non-importers	-0.027**	-0.027***	-0.027***
	(0.011)	(0.010)	(0.010)
Importer differential	0.027***	0.027***	0.027***
	(0.010)	(0.008)	(0.008)
Implied importer effect	0.000	0.000	0.000
Firm fixed effects	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes
Event fixed effects	Yes	Yes	Yes

Notes: The table reports the Hybrid cash-profit-margin results under alternative clustering choices. The specification include firm, year, and event fixed effects unless otherwise indicated. Standard errors are clustered by firm and state. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

The coefficients are unchanged by construction, and the main domestic/non-importer and importer-differential estimates remain statistically significant under each clustering choice. The inference supporting the main result is therefore not specific to the baseline two-way clustering procedure.

Table 19 addresses the influence of extreme cash-profit-margin observations. Profitability ratios can contain large values when sales are unusually low or accounting outcomes are volatile. Re-estimating the model using alternative winsorization thresholds tests whether a small number of observations in the tails of the outcome distribution determine the estimated effects.

Table 19: Winsorization robustness for cash-profit margins

	0.5%	1%	2.5%
Post-treatment window [0,2]			
Domestic/non-importers	-0.036**	-0.033***	-0.024***
	(0.016)	(0.010)	(0.006)
Importer differential	0.039***	0.035***	0.025***
	(0.013)	(0.008)	(0.005)
Implied importer effect	0.003	0.002	0.001
Post-treatment window [0,4]			
Domestic/non-importers	-0.028*	-0.027***	-0.021***
	(0.016)	(0.010)	(0.006)
Importer differential	0.030**	0.027***	0.021***
	(0.013)	(0.008)	(0.005)
Implied importer effect	0.002	0.000	0.000
Firm fixed effects	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes
Event fixed effects	Yes	Yes	Yes

Notes: The table reports cash-profit-margin results using alternative winsorization thresholds for the outcome. The specification include firm, year, and event fixed effects unless otherwise indicated. Standard errors are clustered by firm and state. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

The magnitude of the coefficients decreases somewhat under more aggressive winsorization, as expected when extreme values are compressed, but the qualitative pattern is unchanged. Across all thresholds, the domestic/non-importer effect remains negative, the importer differential remains positive, and the implied importer effect remains close to zero.

9.4.4 Importer-status placebo

The final check evaluates whether the estimation procedure mechanically produces a positive importer differential when importer status contains no systematic information about firm resilience. In Table 20, importer status is randomly permuted 500 times while flood exposure and the estimation-sample structure remain fixed. The preferred specification is re-estimated after each permutation, generating a reference distribution of importer-differential estimates under the imposed placebo assignment mechanism (Ritzwoller et al., 2024). Because importer status is not experimentally assigned, this exercise should be interpreted as a permutation-based diagnostic rather than as an exact randomization test.

Table 20: Importer-status placebo test

Statistic	Value
Actual importer differential	0.035
Actual t-statistic	4.29
Number of placebo assignments	500
Permutation p-value, estimate	0.002
Permutation p-value, t-statistic	0.004
Placebo mean estimate	0.002
Placebo standard deviation	0.006
Placebo 5th percentile	-0.007
Placebo median	0.002
Placebo 95th percentile	0.012

Notes: Importer status is randomly permuted 500 times while flood exposure and the estimation-sample structure remain fixed. The table compares the actual short-run importer differential in cash-profit margins with the resulting placebo distribution. Empirical p-values are computed using the plus-one correction.

The actual short-run importer differential of 0.035 lies well above the placebo distribution, whose mean is 0.002 and whose 95th percentile is 0.012. The corresponding permutation-based p-values are below 0.01. This indicates that a differential as large as the observed estimate is unusual under the imposed placebo reassignment scheme and is unlikely to be generated solely by the mechanics of the estimator.

Taken together, the checks address distinct threats to the baseline interpretation. The cash-profit-margin pattern is stable after accounting for richer pre-shock heterogeneity and state-year shocks, is not driven by a single flood episode or extreme outcome observations, and remains statistically significant under alternative clustering choices. The timing results locate the resilience margin primarily in the immediate post-flood period, while the placebo diagnostic shows that the observed importer differential is unusual under random reassignment of importer status.

9.5 Sector heterogeneity

This subsection examines whether the aggregate cash-profit-margin result is concentrated in a particular manufacturing sector. I re-estimate the preferred Hybrid specification separately for five broad sector groups, allowing both the domestic/non-importer effect and the importer differential to vary across sectors. This exercise helps assess whether the baseline result reflects a broader resilience pattern or is driven by firms operating in a single industry.

The sector-specific estimates should nevertheless be interpreted as descriptive heterogeneity diagnostics. Dividing the sample reduces statistical power, particularly in the smaller sector groups, and differences in statistical significance across columns do not by themselves imply that the underlying coefficients differ significantly across sectors.

Table 21: Sector heterogeneity in cash-profit margin effects: Hybrid estimator

	Heavy industry	Machinery/transport	Agro/resource	Light manuf.	Rubber/plastics/other
Post-treatment window [0,2]					
Domestic/non-importers	-0.028*	-0.075	-0.012	-0.023**	-0.239**
	(0.014)	(0.048)	(0.017)	(0.011)	(0.103)
Importer differential	0.040***	0.064	0.012	0.021	0.263**
	(0.009)	(0.054)	(0.013)	(0.014)	(0.123)
Implied importer effect	0.012	-0.011	0.000	-0.003	0.024
Post-treatment window [0,4]					
Domestic/non-importers	-0.034*	-0.071*	0.015	-0.019*	-0.241***
	(0.019)	(0.039)	(0.017)	(0.010)	(0.046)
Importer differential	0.036***	0.060	-0.005	0.011	0.249***
	(0.013)	(0.045)	(0.012)	(0.010)	(0.060)
Implied importer effect	0.003	-0.010	0.010	-0.008	0.008
Observations	3,775	3,114	2,445	1,556	589
R^2	0.423	0.384	0.441	0.423	0.560
Within R^2	0.006	0.019	0.022	0.012	0.065
Firm fixed effects	Yes	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes	Yes
Event fixed effects	Yes	Yes	Yes	Yes	Yes
Additional controls	No	No	No	No	No
Clustered SE	Firm + State	Firm + State	Firm + State	Firm + State	Firm + State

Notes: The dependent variable is the cash-profit margin. Each column reports estimates from the preferred Hybrid event-study specification for a separate manufacturing sector. The domestic/non-importer effect refers to firms without pre-shock import links; the importer differential is the additional effect for pre-shock importers; and the implied importer effect is the sum of the two coefficients. All specifications include firm, year, and event fixed effects. Standard errors clustered by firm and state are reported in parentheses. Heavy industry includes chemicals and metals; machinery/transport includes machinery and transport equipment; agro/resource includes agro-based and resource-based manufacturing; light manufacturing includes other light manufacturing activities; and rubber/plastics/other includes rubber, plastics, and residual manufacturing industries. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

The sector-specific estimates show that the importer-resilience pattern is not solely driven by one manufacturing sector, although its strength varies across groups. The clearest offsetting pattern appears in heavy industry and in rubber, plastics, and other manufacturing: flood exposure reduces the margins of domestic non-importers, while the importer differential is positive and similar or larger in magnitude. The machinery and transport estimates display the same qualitative pattern, but they are less precisely estimated.

The evidence is weaker in light manufacturing and agro/resource-based industries. In light manufacturing, the domestic/non-importer effect is negative, but the importer differential is smaller and imprecise. In the agro/resource group, neither component displays the baseline offsetting pattern. Across most sectors, however, the implied importer effect remains close to zero.

These results suggest that the aggregate finding is not attributable to a single sector, but they do not imply a uniform resilience effect across manufacturing. In particular, the large coefficients for rubber, plastics, and other manufacturing should be interpreted cautiously because that group contains only 589 observations. Formal tests of coefficient equality would be required to establish whether the importer differential differs statistically across sectors.

9.6 Mechanism diagnostics

The mechanism analysis is intended to identify outcome patterns that are consistent with the channels discussed in the main text, rather than to estimate a formal mediation model. The diagnostics below evaluate candidate outcomes along three dimensions: the direction and magnitude of their post-flood responses, the stability of those responses across post-treatment windows and estimators, and the compatibility of the estimates with pre-treatment dynamics.

Table 22 provides a broad screening exercise across candidate outcomes. The cash-profit margin is included as a benchmark against which the strength of the mechanism evidence can be compared. Table 23 then focuses more closely on labor-cost and distribution-related outcomes, reporting pre-trend diagnostics separately for domestic non-importers and for the importer differential.

Table 22: Mechanism audit across candidate outcomes

Outcome	Estimator	Domestic avg.	Diff. avg.	Importer avg.	Domestic sig.	Diff. sig.	Pre-trend p: diff.
Cash-profit margin	Hybrid	-0.030	0.031	0.001	2/2	2/2	0.341
Cash-profit margin	IPW	-0.025	0.024	-0.001	2/2	2/2	0.141
Cash-profit margin	PS	-0.028	0.027	-0.001	2/2	2/2	0.200
Distribution-to-sales ratio	Hybrid	-0.001	0.001	0.000	0/2	1/2	0.251
Distribution-to-sales ratio	IPW	0.000	0.001	0.001	0/2	2/2	0.605
Distribution-to-sales ratio	PS	0.000	0.001	0.001	0/2	1/2	0.595
Log distribution expenses	Hybrid	-0.052	0.051	-0.002	2/2	2/2	< 0.001
Log distribution expenses	IPW	-0.034	0.041	0.007	2/2	2/2	< 0.001
Log distribution expenses	PS	-0.036	0.043	0.007	2/2	2/2	< 0.001
Log marketing expenses	Hybrid	-0.054	0.062	0.008	2/2	2/2	< 0.001
Log marketing expenses	IPW	-0.037	0.058	0.021	2/2	2/2	< 0.001
Log marketing expenses	PS	-0.040	0.068	0.028	2/2	2/2	< 0.001
Log raw materials	Hybrid	-0.011	0.042	0.030	0/2	0/2	0.151
Log raw materials	IPW	-0.026	0.046	0.019	0/2	0/2	0.135
Log raw materials	PS	-0.025	0.050	0.025	0/2	1/2	0.113
Log sales	Hybrid	-0.020	0.037	0.017	0/2	1/2	0.230
Log sales	IPW	-0.024	0.036	0.012	0/2	1/2	0.108
Log sales	PS	-0.027	0.041	0.014	0/2	1/2	0.176
Log travel expenses	Hybrid	-0.070	0.066	-0.004	2/2	1/2	< 0.001
Log travel expenses	IPW	-0.058	0.063	0.005	0/2	0/2	< 0.001
Log travel expenses	PS	-0.071	0.077	0.006	2/2	0/2	< 0.001
Log wages	Hybrid	-0.015	0.028	0.013	0/2	0/2	< 0.001
Log wages	IPW	-0.022	0.026	0.003	0/2	0/2	< 0.001
Log wages	PS	-0.019	0.030	0.011	0/2	0/2	< 0.001
Wage-to-sales ratio	Hybrid	0.012	-0.017	-0.005	2/2	2/2	0.238
Wage-to-sales ratio	IPW	0.009	-0.014	-0.005	0/2	2/2	0.012
Wage-to-sales ratio	PS	0.013	-0.017	-0.003	2/2	2/2	0.027

Notes: Domestic avg., Diff. avg., and Importer avg. report the arithmetic mean of the estimates over the $[0, 2]$ and $[0, 4]$ post-treatment windows for domestic non-importers, the importer differential, and the implied importer effect, respectively. Domestic sig. and Diff. sig. report the number of windows in which the corresponding estimate is statistically significant at the 5% level. Because the two post-treatment windows overlap, these counts are descriptive and should not be interpreted as independent tests. The final column reports the joint Wald-test p-value for importer-differential leads over $[-5, -2]$. The cash-profit margin is included as a benchmark outcome.

The broad audit distinguishes the stability of the main profitability result from the more mixed evidence on individual channels. The cash-profit-margin pattern is consistent across all three estimators and is supported by non-rejection of the differential pre-trend tests. The wage-to-sales ratio also displays a clear offsetting pattern: flood exposure raises labor-cost intensity among domestic non-importers, while the importer differential is negative. This pattern is strongest in the preferred Hybrid specification, for which the differential pre-trend test does not reject.

Several distribution-related level outcomes display large and statistically significant post-treatment coefficients. In particular, distribution, marketing, and travel expenses decline among domestic non-importers and are partly or fully offset by positive importer differentials. However, the corresponding pre-trend tests reject strongly across estimators. These outcomes therefore provide descriptive evidence that commercial activity evolves differently for importers and non-importers, but they do not offer equally strong support for a causal mechanism interpretation.

The evidence for sales and raw-material use is weaker, with small or imprecisely estimated post-treatment effects. The distribution-to-sales ratio exhibits more favorable pre-treatment diagnostics, but its coefficient magnitudes are small. Overall, the audit identifies labor-cost intensity as the cleanest candidate mechanism, while the evidence concerning distribution-related activity should be interpreted more cautiously.

Table 23 provides a more detailed assessment of the outcomes most closely related to the proposed labor-cost and commercial-capacity channels. In addition to the post-treatment averages and significance counts, it reports separate pre-trend tests for the domestic/non-importer coefficients and the importer-differential coefficients. This distinction is important because a post-flood offsetting pattern is difficult to interpret as evidence of a mechanism when either group was already following a systematically different trajectory before exposure.

Table 23: Labor-cost and distribution diagnostics

Outcome	Estimator	Domestic avg.	Diff. avg.	Importer avg.	Domestic sig.	Diff. sig.	Pre-trend p: domestic	Pre-trend p: diff.
Wage-to-sales ratio	Hybrid	0.012	-0.017	-0.005	2/2	2/2	0.689	0.238
Wage-to-sales ratio	IPW	0.009	-0.014	-0.005	0/2	2/2	0.157	0.012
Wage-to-sales ratio	PS	0.013	-0.017	-0.003	2/2	2/2	0.273	0.027
Log wages	Hybrid	-0.015	0.028	0.013	0/2	0/2	< 0.001	< 0.001
Log wages	IPW	-0.022	0.026	0.003	0/2	0/2	< 0.001	< 0.001
Log wages	PS	-0.019	0.030	0.011	0/2	0/2	< 0.001	< 0.001
Log sales	Hybrid	-0.020	0.037	0.017	0/2	1/2	0.650	0.230
Log sales	IPW	-0.024	0.036	0.012	0/2	1/2	0.482	0.108
Log sales	PS	-0.027	0.041	0.014	0/2	1/2	0.584	0.176
Log distribution expenses	Hybrid	-0.052	0.051	-0.002	2/2	2/2	< 0.001	< 0.001
Log distribution expenses	IPW	-0.034	0.041	0.007	2/2	2/2	0.007	< 0.001
Log distribution expenses	PS	-0.036	0.043	0.007	2/2	2/2	0.001	< 0.001
Distribution-to-sales ratio	Hybrid	-0.001	0.001	0.000	0/2	1/2	0.271	0.251
Distribution-to-sales ratio	IPW	0.000	0.001	0.001	0/2	2/2	0.892	0.605
Distribution-to-sales ratio	PS	0.000	0.001	0.001	0/2	1/2	0.805	0.595
Log marketing expenses	Hybrid	-0.054	0.062	0.008	2/2	2/2	< 0.001	< 0.001
Log marketing expenses	IPW	-0.037	0.058	0.021	2/2	2/2	< 0.001	< 0.001
Log marketing expenses	PS	-0.040	0.068	0.028	2/2	2/2	< 0.001	< 0.001
Log travel expenses	Hybrid	-0.070	0.066	-0.004	2/2	1/2	< 0.001	< 0.001
Log travel expenses	IPW	-0.058	0.063	0.005	0/2	0/2	< 0.001	< 0.001
Log travel expenses	PS	-0.071	0.077	0.006	2/2	0/2	< 0.001	< 0.001

Notes: Domestic avg., Diff. avg., and Importer avg. report averages over the overlapping $[0, 2]$ and $[0, 4]$ post-treatment windows for domestic non-importers, the importer differential, and the implied importer effect, respectively. Domestic sig. and Diff. sig. report the number of windows in which the corresponding estimate is statistically significant at the 5% level. The final two columns report joint Wald-test p-values for the domestic/non-importer and importer-differential leads over $[-5, -2]$. Because the post-treatment windows overlap, the significance counts are descriptive rather than independent tests.

The focused diagnostics indicate that the wage-to-sales ratio provides the strongest support for the labor-cost channel in the preferred Hybrid specification. Its domestic/non-importer and importer-differential estimates move in opposite directions, and neither set of pre-treatment coefficients rejects the null of no pre-trends. The corresponding PS and IPW estimates display a similar post-treatment pattern, although their importer-differential pre-trend tests are weaker.

The evidence for distribution capacity is more mixed. Distribution and marketing expenses display strong offsetting post-treatment coefficients across estimators, but their pre-treatment tests reject sharply. By contrast, the distribution-to-sales ratio has more favorable pre-treatment diagnostics but much smaller estimated effects. Log sales also displays acceptable pre-treatment dynamics, although its post-treatment effects are generally imprecise.

These results support a cautious interpretation. The evidence is consistent with importers experiencing less labor-cost pressure after flood exposure, while the distribution-related outcomes provide suggestive evidence that importers preserve commercial activity more effectively. Because several candidate mechanism outcomes exhibit differential pre-trends, the mechanism estimates should not be interpreted as formal causal mediation effects.